arXiv:1604.07130v1 [math.OC] 25 Apr 2016

Worst-case Complexity of Cyclic Coordinate Descent: $O(n^2)$ Gap with Randomized Version

Ruoyu Sun * Yinyu Ye†

Apr 24, 2016

Abstract

This paper concerns the worst-case complexity of Gauss-Seidel method for solving a positive semi-definite linear system; or equivalently, that of cyclic coordinate descent (C-CD) for minimizing a convex quadratic function. The known provable complexity of C-CD can be $O(n)$ times slower than gradient descent (GD) and $O(n^2)$ times slower than randomized coordinate descent (R-CD). However, these gaps seem rather puzzling since so far they have not been observed in practice; in fact, C-CD usually converges much faster than GD and sometimes comparable to R-CD. Thus some researchers believe the gaps are due to the weakness of the proof, but not that of the C-CD algorithm itself.

In this paper we show that the gaps indeed exist. We prove that there exists an example for which C-CD takes at least $\tilde{O}(n^3\kappa)$ or $\tilde{O}(n^4\kappa_{\mathrm{CD}})$ operations, where κ is the condition number, κ_{CD} is a well-studied quantity that determines the convergence rate of R-CD, and $\tilde{O}(\cdot)$ hides the dependency on $\log 1/\epsilon$. This implies that C-CD can indeed be $O(n)$ times slower than GD, and $O(n^2)$ times slower than R-CD in the worst case. Our result establishes one of the few examples in continuous optimization that demonstrates a large gap between a deterministic algorithm and its randomized counterpart. Based on the example, we establish several almost tight complexity bounds of C-CD for quadratic problems. One difficulty with the analysis of the constructed example is that the spectral radius of a non-symmetric iteration matrix does not necessarily constitutes a lower bound for the convergence rate.

1 Introduction

Coordinate descent (CD) algorithms have been very popular recently for solving large-scale unconstrained optimization problems. In the most basic form, cyclic CD (C-CD) optimizes over one variable at a time with other variables fixed, and the variables are chosen according to a fixed order. C-CD for minimizing convex quadratic functions is equivalent to the Gauss-Seidal method for solving a PSD (Positive Semi-Definite) linear system. A landmark in the history of CD algorithms is the establishment of the convergence rate of randomized CD (R-CD) [1, 2], a variant which updates variables randomly. In particular, R-CD was shown to be $\mathcal{O}(1)$ to $\mathcal{O}(n)$ times faster than GD, where n is the number of variables. There are some recent efforts in understanding the convergence speed of C-CD [3–5], but the best existing rates can be $\mathcal{O}(1)$ to $\mathcal{O}(n)$ times worse than GD (gradient descent). In practice, C-CD was found to converge much faster than GD, and sometimes even faster than R-CD. The huge gap between the existing results and numerical experiments makes the true convergence rate of C-CD a puzzle.

*Department of Management Science and Engineering, Stanford University, Stanford, CA. ruoyus@stanford.edu.
†Department of Management Science and Engineering, Stanford University, Stanford, CA. yyye@stanford.edu.

1.1 Background

Gauss-Seidel (G-S) method, first proposed by Gauss in 1823, is one of the oldest iterative algorithms. It can be used to solve any system of linear equations, and the convergence was guaranteed when the coefficient matrix is strictly diagonally dominant or PSD. Regarding the convergence speed, it is well-known that for some special matrices, asymptotically G-S method converges twice as fast as Jacobi method, a parallel version of G-S method (see, e.g., [6]).

Coordinate descent methods also enjoy a long history in nonlinear optimization. As the most basic version, cyclic CD for minimizing convex quadratic functions is equivalent to Gauss-Seidel method for solving PSD (Positive Semi-Definite) linear systems. There are many extensions/variants of this basic form, using different update orders, different block size, different subproblems, etc. Due to the simplicity, CD-type methods are one of the most widely used class of methods in practice (see, e.g., [7–19]). Perhaps for the same reason CD-type methods were not extensively studied by optimization researchers before, say, 2010. Most early works focused on the exact conditions for the convergence or the quality of convergence, not the quantity of convergence rate (see, e.g., [20–23]).

Recently, there has been a surge of interest for CD-type methods due to their efficiency for solving large-scale optimization problems. The efficiency was not only observed in practice, but also validated by theory: randomized BCGD (Block Coordinate Gradient Descent) is shown to converge $\frac{nL}{\sum_i L_i}$ (with maximal value n) times faster than GD in solving smooth convex problems [1, 2], where n is the number of blocks, L is the Lipschitz constant of the gradient, and L_i is the per-block Lipschitz constant of the gradient. Note that the introduction of randomized update order is crucial since it makes the analysis of CD-type methods quite simple. Ever since then, randomized update rule has been a new standard for theoretical analysis of CD-type methods [24–34]. For example, the results on the convergence speed of randomized versions have been extended to composite objective functions [25, 26]. Rather surprisingly, Lee and Sidford [35] showed that accelerated R-CD proposed by Nesterov [2] can be efficiently implemented so that its overall complexity is $\mathcal{O}(\frac{nL}{\sum_i L_i})$ times better than conjugate gradient method for solving PD (positive definite) linear systems. Allen-Zhu and Orecchia [36] adopted accelerated R-CD (together with other ideas) to design faster algorithms for solving packing and covering LP (Linear Programming).

With all these nice theoretical results on randomized CD methods, one may wonder whether the same results can be achieved for cyclic CD methods. There are several reasons for studying cyclic CD methods. First, in practice, it is not always easy to randomly pick coordinates. A computer system usually divides the whole set of variables in the hard disk into many chunks and load one chunk of variables into the memory. As the typical computer architecture consists of multiple layers of caches, the above "cut and load" process may happen several times. If the transition between the cache, the memory and the hard disk is fast enough, applying randomized coordinate selection can be feasible; otherwise it might be beneficial to process all coordinates one by one. Therefore cyclic CD methods might be more hardware-friendly than randomized versions in some scenarios.

Second, the complexity of deterministic algorithms is theoretically important (partly because generating fully random bits is a highly non-trivial task). For example, for testing the primality of integers, the first deterministic algorithm that runs in polynomial time was regarded as a great achievement of the complexity theory [37], even though polynomial time randomized algorithms had been known for decades. In optimization, it is an interesting open question whether there exists a version of deterministic simplex method that can solve LP in polynomial time. In our case, as randomized CD methods are the crucial components of the fastest (randomized) algorithms for several class of problems mentioned above (PD linear system and positive LP), it is an interesting question whether the same complexity can be achieved by

deterministic algorithms. Understanding the complexity of cyclic CD will serve as an important first step towards answering this question.

There have been some recent efforts to understand the convergence speed of cyclic BCGD (Block Coordinate Gradient Descent) [3–5, 38, 39]. In BCGD, a typical stepsize for the i-th block is $1/L_i$, and for randomized BCGD this stepsize can lead to an optimal rate. For minimizing convex quadratic functions, cyclic CGD with stepsize $1/L_i$ is just cyclic CD (i.e. exactly minimizing each subproblem), and also equivalent to Gauss-Seidel method. The best known convergence rate of cyclic CGD with stepsize $1/P_i$ [5, 40] for minimizing quadratic functions [1] is $\mathcal{O}(\frac{L}{\min_i P_i}\log^2 n)$ times worse than GD when $P_i \leq L$. This result has two implications.

- For simplicity, consider the equal-diagonal case (i.e. $L_i = L_1,\ \forall i$). The best known convergence rate of C-CD or Gauss-Seidal method is $\mathcal{O}(\frac{L}{L_1}\log^2 n)$ times *worse* than GD. In contrast, the rate of R-CD is $\mathcal{O}(\frac{L}{L_1})$ times *better* than GD. Note that the gap L/L_1 can be as large as n, and for the Wishart random matrix $A = U^T U$ where U has i.i.d. Gaussian entries $L/L_1 = \mathcal{O}(\sqrt{n})$, thus the gap L/L_1 indeed can be very large. Such opposite results of C-CD and R-CD look rather strange, and do not match the numerical experiments. In fact, almost all simulation results show that C-CD converges much faster than GD and sometimes faster than randomized CD. The existing results seem so weak that they even make a wrong prediction on whether C-CD is faster than GD.

- The convergence rate of cyclic CGD with small stepsize $1/L$ is only $\mathcal{O}(\log^2 n)$ times worse than GD with stepsize $1/L$, thus much better than cyclic CGD with large stepsize $1/L_i$ (or equivalently, C-CD). The former comparison is somewhat acceptable since cyclic CGD now takes the same stepsize as GD, but the latter comparison shows a "stepsize dilemma": in the range $[1/L, 1/L_i]$, smaller stepsize leads to faster convergence of cyclic CGD by known theory, but in practice leads to slower convergence. Note that randomized CGD can be much faster than GD mainly because it can take large stepsize $1/L_i$, and its maximal allowable stepsize $2/L_i$ is much larger than the maximal allowable stepsize of GD $2/L$. Meanwhile, the maximal allowable stepsize for cyclic CGD is still $2/L_i$ but this does not help cyclic CGD. It is a common rule of thumb to pick aggressive stepsizes to achieve faster convergence; now this rule seems to fail for cyclic CGD by known theory.

It is very tempting to think that we might be able to prove C-CD is faster than GD, or even comparable to randomized CD. The huge gap between the theory and the practice as well as the stepsize dilemma might just be because of the weakness of the proof techniques. This impression may be enhanced when we reflect on the existing proofs of upper bounds. The proof idea of [3, 5] is to view cyclic BCGD as an inexact version of GD, and the major effort is spent on bounding the difference between cyclic BCGD and GD. One obvious drawback of such a proof framework is that it cannot show a better convergence rate than GD; even if the difference is zero, only the same rate will be established. This raises a question related to deeper understanding of CD-type methods: should we really view CD-type methods as inexact GD? The fact that randomized CD methods are faster than GD implies that randomized versions should *not* be viewed as inexact GD. It seems possible that there exists a different proof framework for C-CD that leads to better convergence rates.

To summarize, the known theory explains neither why C-CD is faster than GD in practice, nor why larger stepsize within range $[1/L, 1/L_i]$ leads to faster convergence of cyclic CGD in practice. In addition, the existing proof framework seems to be not strong enough to provide the tightest bound. However, we will show that the known results are almost the best we can hope for.

[1] [5] studied possibly non-strongly convex functions and established sublinear convergence rates, but the same techniques can lead to the proof of similar linear convergence results for strongly convex functions.

1.2 Summary of Contributions

We will focus on the worst-case complexity of C-CD for minimizing convex quadratic functions $\min_{x\in\mathbb{R}^n} \frac{1}{2}x^TAx - b^Tx$ (i.e. Gauss-Seidel method for solving a PSD linear system $Ax = b$). It is not hard to extend our results to cyclc BCD (the block size can be larger than 1) and cyclic BCGD with different stepsizes, but for simplicity we will not discuss these extensions. In the following, we will say an algorithm has complexity $\tilde{\mathcal{O}}(g(n,\theta))$, if it takes $\mathcal{O}(g(n,\theta)\log(1/\epsilon))$ unit operations to achieve relative error $\frac{f(x)-f^*}{f(x^0)-f^*} \leq \epsilon$. For instance, it is known that GD has complexity $\tilde{\mathcal{O}}(n^2\kappa)$, where $\kappa = \lambda_{\max}(A)/\lambda_{\min}(A)$ is the condition number, and R-CD has complexity $\tilde{\mathcal{O}}(n^2\kappa_{CD})$, where $\kappa_{CD} = \sum_i A_{ii}/(n\lambda_{\min}(A))$.

For the upper bound, we strengthen the existing proof to give several new bounds that depend on different parameters of the problem. Our main contribution is to establish several lower bounds by analyzing the convergence rate of a simple class of examples. Our discovery is that the upper bounds are "almost" tight (sometimes up to $\mathcal{O}(\log^2 n)$ factor) in the equal-diagonal case (i.e. all diagonal entries $L_i = L_1, \forall i$). We summarize our findings for the worst-complexity in the *equal-diagonal case* below (the non-equal-diagonal case is quite subtle and related to conjectures on Jacobi-preconditioning; see Section 3.2). We will discuss two types of bounds, one does not involve L/L_1 (recall $L = \lambda_{\max}(A), L_1 = A_{11}$) and the other does.

- In terms of κ or κ_{CD}, the worst-case complexity of C-CD is

$$\tilde{\mathcal{O}}(n^3\kappa) \text{ or } \tilde{\mathcal{O}}(n^4\kappa_{CD}). \tag{1}$$

  Both bounds are tight as for our example both bounds are achieved. This implies that C-CD can be $\mathcal{O}(n)$ times slower than GD and $\mathcal{O}(n^2)$ times slower than R-CD.

- It is more precise to characterize the complexity using an extra parameter $\tau = L/L_1$ together with κ or κ_{CD}. The lower bound for the complexity of C-CD is

$$\tilde{\mathcal{O}}\left(n^2\kappa\tau\right) \text{ or } \tilde{\mathcal{O}}\left(n^2\kappa_{CD}\tau^2\right), \tag{2}$$

  which is τ times worse than GD or τ^2 times worse than R-CD. The range of τ is $[1, n]$, and for Wishart random matrices $\tau = \mathcal{O}(\sqrt{n})$, thus the gap τ or τ^2 is indeed very large. There is an $\mathcal{O}(\log^2 n)$ gap between these lower bounds and the corresponding upper bounds, but the $\log^2(n)$ factor was previously known to be unavoidable in some cases. See more detailed discussion in Section 3.1.

To prove these lower bounds, we only need to estimate the convergence rate of our specific examples. There are at least two difficulties. Firstly, there is no closed form expression of the spectral radius of the iteration matrix and we need to consider the limiting behavior of a class of examples (still with fixed n). Secondly and more interestingly, the spectral radius does not directly lead to a lower bound of the convergence rate when the iteration matrix is *non-symmetric*, and we need to explore some special structure of the examples.

We also perform some numerical experiments to see whether the practical performance of C-CD matches our theory. We summarize our numerical findings below (again, only consider the equal-diagonal case).

- One empirical implication of our theory is that C-CD converges slowly when L/L_1 is large or the off-diagonal entries are large. Interestingly, this phenomenon is indeed observed in our numerical experiments: C-CD converges faster than R-CD when off-diagonal entries are small but slower than R-CD when off-diagonal entries are large. This phenomenon is rather universal as it is observed for a wide variety of generative models of the coefficient matrix.

- In almost all scenarios (except random perturbations of our example) C-CD converges much *faster* than GD, which is *opposite* to the theory which says that C-CD converges $\mathcal{O}(1)$ to $\mathcal{O}(n)$ times *slower* than GD. The theory also says that C-CD converges up to $\mathcal{O}(n^2)$ times slower than R-CD, but the gap between C-CD and R-CD is very small compared to n. Therefore, our experiments show that C-CD is still a very competitive algorithm even though we prove it converges slowly in the worst case.

The huge gap between the practice and the theory partially explains why the $\mathcal{O}(n^2)$ gap has never been reported. It reveals the weakness of the worst-case analysis of convergence speed, which has been universal in optimization papers; nevertheless, this weakness has long been recognized and is considered very difficult to tackle.

1.3 Implications

We summarize our discoveries and discuss some interesting open-ended questions motivated by this paper.

Gap Between Deterministic and Randomized Algorithms. We discover for the first time that C-CD (and Gauss-Seidel method) can be $\mathcal{O}(n^2)$ times slower than their randomized counterparts. This is one of the few examples in continuous optimization that a large gap between a deterministic algorithm and its randomized counterpart is established.

Stepsize Dilemma. We find that the "stepsize dilemma" does exist both in theory and in certain real examples. More specifically, in our example C-CD has an allowable stepsize range $[0, 2)$ which is much larger than the allowable stepsize range of GD $[0, 2/n)$, but large stepsizes do not help C-CD. In fact, in the region $[1/n, 1]$ increasing the stepsize of C-CD would rather decrease the convergence speed for our example. Again for most practical examples we do not observe such a stepsize dilemma.

Fundamental Gap Between Deterministic and Randomized CD? Our results only establish a large gap between a single deterministic version of CD and R-CD. A natural question arises: is there a fundamental gap between deterministic CD and randomized CD? There has to be an answer: either we can prove a large lower bound for *all* deterministic CD methods, or we can find one deterministic CD that performs close to R-CD. Both possibilities are very interesting. For the latter possibility, there are a few candidates such as CGD with stepsize other than $1/L_i$ (equivalent to SOR, i.e. successive over relaxation) and double sweep method (a.k.a. symmetric SOR), but they are far worse than R-CD for our example.

Deterministic Complexity. Recent progress on the complexity of some important classes of problems (e.g. PD linear systems, positive LP) is based on randomized versions of CD methods. As we have established a large gap between C-CD and R-CD, it is unclear whether the same complexity can be achieved for deterministic algorithms. CG (conjugate gradient) is still the fastest *deterministic* iterative algorithm for solving PSD linear systems, even though accelerated R-CD is faster in a probabilistic sense.

New Type of Analysis? It is an interesting question how to explain the large gap between the theory and the practical performance of C-CD. This kind of gaps usually lead to novel theoretical advances. One famous example is the smoothed analysis developed by Spielman and Teng [41] that aims to explain such a gap for the simplex method. What type of analysis is suitable for explaining the practical performance of C-CD (e.g. why is it usually much faster than GD)?

Related Algorithms. We hope this research will shed light on the study of related algorithms, such as Kaczmarz method, SGD (Stochastic Gradient Descent) and ADMM (Alternating Direction of Multiplier Method). Since Kaczmarz method can be interpreted as coordinate descent when solving consistent overdetermined linear systems, our results also imply that classical cyclic Kaczmarz method proposed in 1937 [42] can be $\mathcal{O}(n^2)$ times slower than its randomized counterpart proposed recently [43]. As for ADMM, it was

recently found that the cyclic version with at least 3 blocks can be divergent [44], while randomly permuted version converges in expectation for solving linear systems [45], so a fundamental gap between cyclic versions and randomly permuted versions exsits. Nevertheless, it was also known that for certain problems (e.g. strongly convex) the small-stepsize versions of cyclic ADMM can be convergent [46–48]. Based on our results, we conjecture that in these cases cyclic ADMM achieve worse complexity than ALM (Augmented Lagrangian Method, the "batch version" of ADMM) and randomized versions of ADMM.

1.4 Notations and Organization

Most notations in this paper are standard. Throughout the paper, $A \in \mathbb{R}^{n\times n}$ is a symmetric positive semi-definite matrix. Let $L = \lambda_{\max}(A), \lambda_{\min}(A)$ denote the maximum eigenvalue and minimum *non-zero* eigenvalue of A respectively; sometimes we omit the argument A and just use $\lambda_{\max}$ and $\lambda_{\min}$. The condition number of A is defined as $\kappa = \frac{\lambda_{\max}(A)}{\lambda_{\min}(A)}$. Denote A_{ij} as the (i,j)-th entry of A and $L_i = A_{ii}$ as the i-th diagonal entry of A. We use redundant notations L and L_i to be consistent with the optimization literature: L represents the global Lipschitz constant and L_i represents the i-th coordinate Lipschitz constant of the gradient of the function $\frac{1}{2}x^TAx$. We denote $\mathcal{R}(A) = \{Ax \mid x \in \mathbb{R}^n\}$ as the range space of A. Denote $A^\dagger$ as the pseudo-inverse of A, which can be defined as $V\text{diag}\{1/\lambda_1, \ldots, 1/\lambda_r, 0, \ldots, 0\}V^{-1}$ when the eigen-decomposition of A is $V\text{diag}\{\lambda_1, \ldots, \lambda_r, 0, \ldots, 0\}V^{-1}$, where $\lambda_1, \ldots, \lambda_r$ are all the non-zero eigenvalues of A.

The less widely used (or less intuitive) notations are summarized below. We denote $L_{\max} = \max_i L_i$ and $L_{\min} = \min_i L_i$ as the maximum/minimum per-coordinate Lipschitz constant (i.e. maixmum/minimum diagonal entry of A), and $L_{\text{avg}} = (\sum_{i=1}^n L_i)/n$ as the average of the diagonal entries of A (which is also the average of the eigenvalues of A). Denote $\kappa_{CD} = \frac{L_{\text{avg}}}{\lambda_{\min}(A)}$ which is a well-studied quantity that characterizes the convergence rate of R-CD. We usually use Γ to denote the lower triangular part of matrix A with diagonal entries, i.e. $\Gamma_{ij} = A_{ij}$ iff $i \le j$. We also use D_A to denote the diagonal matrix consisting of diagonal entries of A. Sometimes we use the notation $\tau = \frac{L}{L_{\min}}$, a crucial ratio that characterizes the difference between GD, C-CD and R-CD.

The rest of the paper is organized as follows. In Section 2, we review the algorithms discussed in the paper. In Section 3, we present our theoretical results on the complexity of C-CD as well as the comparison of C-CD with other algorithms. Section 4 is devoted to the proof of the main result Theorem 3.1. The main ideas and the outline of the proof will be provided before the formal proof. In Section 5, we present some numerical experiments. The proofs of results other than Theorem 3.1 are provided in the appendix.

2 Preliminaries: Several Algorithms

In this section we will review GD and several variants of CD. For simplicity, we present the versions for minimizing convex quadratic functions instead of the most general form. More specifically, we consider solving

$$\min_{x\in\mathbb{R}^n} \frac{1}{2}x^TAx - b^Tx,$$

where $A \in \mathbb{R}^{n\times n}$ is a symmetric PSD (positive semi-definite) matrix, $b \in \mathcal{R}(A)$ and $A_{ii} \neq 0,\ \forall i$. All the optimal solutions of the problem satisfy the system of linear equations

$$Ax = b.$$

When A is non-singular (thus positive definite), the unique minimizer $x = A^{-1}b$ is the unique solution to the linear system. When A is singular, there are infinitely many optimal solutions of the form $A^{\dagger}b + x^{\perp}$ where $Ax^{\perp} = 0$.

GD (gradient descent) is one of the most basic iterative algorithms. Starting at $x^0 \in \mathbb{R}^n$, GD proceeds as follows:

$$x^{k+1} = x^k - \frac{1}{L}\nabla f(x) = x^k - \frac{1}{L}(Ax - b).$$

There are many other choices of stepsizes instead of $1/L$, but we choose this stepsize because it is simple and already leads to the standard complexity $\mathcal{O}(n^2\kappa)$.

The C-CD algorithm starts at $x^0 = x^{0,0}$ and proceeds as follows:

$$\begin{aligned} &x^{k,j} = x^{k,j-1} - \frac{A(j,:)x^{k,j-1}}{A_{jj}}e_j, \quad j = 1, 2, \dots, n, \\ &x^{k+1} = x^{k,n},\ x^{k+1,0} = x^{k+1}. \end{aligned} \tag{3}$$

We can write the above update equation as a simple matrix recursion

$$x^{k+1} = (I - \Gamma^{-1}A)x^k,$$

where Γ is the lower triangular part of A with diagonal entries, i.e., $\Gamma_{ij} = A_{ij}, 1 \le j \le i \le n$. We denote the update matrix as

$$M = I - \Gamma^{-1}A.$$

R-CD (randomized coordinate descent) algorithm starts at z^0 and proceeds as follows:

$$\begin{aligned} &\text{FOR } j = 1, 2, \dots \\ &\quad \text{Randomly pick } t \in \{1, \dots, n\} \text{ uniformly at random,} \\ &\quad z^{j+1} = z^j - \frac{A(t,:)z^j}{A(t,t)}e_t. \end{aligned} \tag{4}$$

The output of R-CD is a sequence $(z^1, z^2, \dots)$. We further define

$$x^k = z^{kn}, \quad k = 0, 1, 2, \dots$$

to be comparable with GD and C-CD. Here k can be viewed as the index of "epoch", where each epoch consists of n iterations.

We also consider RP-CD (randomly permuted coordinate descent). Starting from $x^0 = x^{0,0}$, the algorithm proceeds as follows.

$$\begin{aligned} &\text{At epoch } k, \text{ pick a permutation } \sigma_k \text{ uniformly at random from the set of all permutations.} \\ &\text{FOR } j = 1, 2, \dots, n, \\ &\quad t = \sigma_k(j); \\ &\quad x^{k,j} = x^{k,j-1} - \frac{A(t,:)x^{k,j-1}}{A_{tt}}e_t; \\ &\quad x^{k+1} = x^{k,n},\ x^{k+1,0} = x^{k+1}. \end{aligned} \tag{5}$$

According to [45, Section II.A], the recursion formula of RP-CD is

$$x^k = (I - \Gamma_{\sigma_k}^{-1}A)x^{k-1}, \tag{6}$$

where σ_k is the permutation used in the k-th iteration, and $L_\sigma \in \mathbb{R}^{n\times n}$ is defined by

$$\Gamma_\sigma(\sigma(i),\sigma(j)) \triangleq \begin{cases} A_{\sigma(i),\sigma(j)} & j \le i. \\ 0 & j > i, \end{cases} \tag{7}$$

For example, when $n = 3$ and $\sigma = (231)$,

$$\Gamma_{(231)} = \begin{bmatrix} A_{11} & A_{12} & A_{13} \\ 0 & A_{22} & 0 \\ 0 & A_{32} & A_{33} \end{bmatrix}.$$

3 Worst-case Complexity of Cyclic Coordinate Descent

We first prove a result on the convergence rate of C-CD for solving quadratic problems, which leads to upper bounds of the complexity. The matrix A is allowed to be positive semi-definite. We assume $b \in \mathcal{R}(A)$ since otherwise the minimum value of $\min_x x^T Ax - 2b^T x$ will be $-\infty$. Without loss of generality, we assume $A_{ii} \neq 0,\ \forall i$, since when some $A_{ii} = 0$ all entries in the i-th row and the i-th column of A should be zero, which means that the i-th variable does not affect the objective and thus can be deleted. The proof of this result will be given in Appendix A.1.

Proposition 3.1. *(Upper bound of C-CD) Consider the quadratic minimization problem* $\min_x f(x) \triangleq x^T Ax - 2b^T x$ *where* $A \in \mathbb{R}^{n\times n}$ *is positive semi-definite,* $b \in \mathcal{R}(A)$ *and* $A_{ii} \neq 0,\ \forall i$. *For any* $x^0 \in R^n$, *let* x^k *denotes the output of C-CD after* k *cycles, then*

$$f(x^{k+1}) - f^* \le \min\left\{1 - \frac{1}{n\kappa}\frac{L_{\min}}{L_{\text{avg}}}, 1 - \frac{L_{\min}}{L(2+\log n/\pi)^2}\frac{1}{\kappa}\right\}(f(x^k) - f^*). \tag{8a}$$

$$f(x^{k+1}) - f^* \le \min\left\{1 - \frac{1}{n^2\kappa_{CD}}\frac{L_{\min}}{L_{\text{avg}}}, 1 - \frac{L_{\min}L_{\text{avg}}}{L^2(2+\log n/\pi)^2}\frac{1}{\kappa_{CD}}\right\}(f(x^k) - f^*). \tag{8b}$$

Here, f^* *is the minimum value of the function* f, $\kappa = \frac{\lambda_{\max}(A)}{\lambda_{\min}(A)}$ *is the condition number in which* $\lambda_{\min}(A)$ *denotes the minimum non-zero eigenvalue of* A, $L = \lambda_{\max}(A)$, $L_{avg} = \sum_i A_{ii}/n$, $L_{\min} = \min_i A_{ii}$ *and* $\kappa_{CD} = \frac{L_{\text{avg}}}{\lambda_{\min}(A)}$.

As mentioned in the introduction, we will say an algorithm has complexity $\tilde{\mathcal{O}}(g(n,\theta))$, if it takes $\mathcal{O}(g(n,\theta)\log(1/\epsilon))$ unit operations to achieve relative error ϵ. Using the fact $-\ln(1-z) \ge -z, z \in (0,1)$ one can immediately show that to achieve $(1-1/u)^k \le \epsilon$ one only needs $k \ge u\log(1/\epsilon)$ iterations.

Consider the equal-diagonal case (i.e. $L_i = L_1,\ \forall i$) for now and we will discuss the general case later in Section 3.2. Since each cycle takes $\mathcal{O}(n^2)$ operations, Proposition 3.1 implies that C-CD has complexity at most

$$\min\left\{\tilde{\mathcal{O}}\left(n^3\kappa\right), \tilde{\mathcal{O}}\left(n^2\kappa\log^2 n\frac{L}{L_1}\right)\right\}. \tag{9}$$

As GD has complexity $\mathcal{O}(n^2\kappa)$, the first bound in (9) is $\mathcal{O}(n)$ times worse than GD, and the second bound is $\mathcal{O}(\frac{L}{L_1}\log^2 n)$ times worse than GD. We will compare these two types of bounds later. The inequality 8b implies the following upper bounds on the complexity of C-CD:

$$\min\left\{\tilde{\mathcal{O}}\left(n^4\kappa_{CD}\right), \tilde{\mathcal{O}}\left(n^2\kappa_{CD}\log^2 n\left(\frac{L}{L_1}\right)^2\right)\right\}. \tag{10}$$

As R-CD has complexity $\mathcal{O}(n^2\kappa_{CD})$ (in probabilistic sense), the first bound in (10) is $\mathcal{O}(n^2)$ times worse than R-CD, and the second bound is $\mathcal{O}((\frac{L}{L_1})^2\log^2 n)$ times worse than R-CD.

The second type of upper bound which depends on L/L_1 was studied before. For example, [40] studied the SOR method and recently [5] studied general convex problems. Both works obtained a bound that is $\mathcal{O}((\frac{L}{L_1})^2\log^2 n)$ times worse than R-CD and both used a classical result that the norm of the triangular truncation operator is $\mathcal{O}(\log n)$.

Our result Proposition 3.1 provides a new type of bound that only depends on κ or κ_{CD} in the equal-diagonal case. We give a unified proof framework that leads to both types of bounds. Our proof can be divided into two stages. The first stage is to relate the convergence rate with the spectral norm of a matrix $\Gamma^{-1}A\Gamma^{-T}$, which can be proved by two different approaches (from different perspectives): one is from optimization which views C-CD as inexact GD; the other is from linear algebra which studies the spectral radius of the iteration matrix $I-\Gamma^{-1}A$. Note that $\Gamma^{-1}A$ is non-symmetric, thus the latter method requires an extra symmetrization technique which relaxes the spectral radius by the spectral norm. As we will see later, such a technique cannot be used in the proof of the lower bound, and other techniques are needed for that proof. In the second stage, we estimate $\|\Gamma^{-1}A\Gamma^{-T}\|$ via two different methods, leading to the two types of bounds. As discussed later in Section 3.1, each bound is tight in one scenario, thus the two bounds cannot be combined into one single bound.

It is very tempting to think that the huge gap between C-CD and GD/R-CD is just due to the weakness of the proof techniques, since people have not observed (at least have not reported to our knowledge) such a huge gap in practice. In particular, it seems very strange that C-CD converges much slower than GD, since people expect that it should be faster than GD. Rather surprisingly, we will show that the bounds given in Proposition 3.1 is actually tight when all diagonal entries are the same (possibly up to $\log^2 n$ factor).

Theorem 3.1. *(Lower bound of CD) For any $x^0 \in R^n$, any $\delta \in (0,1]$, there exists a quadratic function $f(x) = x^TAx + 2b^Tx$ such that*

$$f(x^k) - f^* \geq (1-\delta)\left(1 - \frac{2\pi^2}{n\kappa}\right)^{2k+2}(f(x^0) - f^*) \tag{11a}$$

$$f(x^k) - f^* \geq (1-\delta)\left(1 - \frac{2\pi^2}{n^2\kappa_{CD}}\right)^{2k+2}(f(x^0) - f^*) \tag{11b}$$

where x^k denotes the output of C-CD after k cycles, f^* is the minimum of the objective function f, $\kappa = \frac{\lambda_{\max}(A)}{\lambda_{\min}(A)}$ is the condition number, $\kappa_{CD} = \frac{L_{\text{avg}}}{\lambda_{\min}(A)} = \frac{\sum_i A_{ii}}{n\lambda_{\min}(A)}$.

The proof of this result will be given in Section 4. The example we construct is simple: all diagonal entries of A are 1 and all off-diagonal entries are c, where c is a constant close to 1. It is known that the SDD (symmetric diagonally dominant) system can be solved in almost linear time (see, e.g., [49–51] and the references therein). While for SDD system the off-diagonal entries have very small magnitude, the hard instance we construct can be viewed as the opposite of the SDD system: the off-diagonal entries are chosen as large as possible so that the matrix remains positive definite.

Theorem 3.1 implies that the complexity of C-CD for our example is at least

$$\tilde{\mathcal{O}}\left(n^3\kappa\right) \text{ or } \tilde{\mathcal{O}}\left(n^4\kappa_{CD}\right). \tag{12}$$

These lower bounds match the first upper bounds given in (9) and (10). This means that the $\mathcal{O}(n)$ gap between C-CD and GD, and the $\mathcal{O}(n^2)$ gap between C-CD and R-CD indeed exist! For our example the ratio $\tau = \frac{L}{L_1} = \frac{nL}{\sum_i L_i} \leq n$, thus we immediately obtain another lower bound

$$\tilde{\mathcal{O}}\left(n^2\kappa\tau\right) \text{ or } \tilde{\mathcal{O}}\left(n^2\kappa_{CD}\tau^2\right). \tag{13}$$

These lower bounds "almost" match the second upper bounds given in (9) and (10) except for an extra $\mathcal{O}(\log^2 n)$ factor. The bounds (13) are more precise than the bounds (12). Though we have claimed that C-CD can be $\mathcal{O}(n)$ times slower than GD and $\mathcal{O}(n^2)$ times slower than R-CD, a more precise statement on the *worst-case* complexity of the three methods is the following:

- C-CD is roughly $\mathcal{O}(\tau)$ times slower than GD, and R-CD is $\mathcal{O}(\tau)$ times faster than GD. When τ achieves the maximum $\mathcal{O}(n)$, C-CD is $\mathcal{O}(n)$ times slower than GD and R-CD is $\mathcal{O}(n)$ times faster than GD, which implies C-CD is $O(n^2)$ times slower than R-CD.

Note that in the above statement "method 1 is X-times slower than method 2" does not mean that method 1 is always slower than method 2 (of course rarely can one make such a strong statement), but that "the worst-case complexity of method 1 is X-times worse than that of method 2, and both complexity bounds can be simultaneously achieved". We say "roughly" when comparing C-CD with GD because of the $\log^2 n$ gap between the upper bound and the lower bound; see more detailed discussion on this gap in the following subsection.

3.1 Necessity of Two Types of Bounds

Consider the equal-diagonal case (i.e. $L_i = L_1,\ \forall i$) in this subsection. We will explain the relation between the two types of bounds, one does not involve $\tau = L/L_1$ and another does. We argue that it is not easy, if not impossible, to obtain one single tight bound. We will also suggest slightly stronger bounds that might be the tightest based on the current parameters (again, for the equal-diagonal case).

We denote two bounds related to κ as $B_1 = \mathcal{O}(n^3\kappa)$ and $B_2 = \mathcal{O}(n^2\kappa\tau\log^2 n)$; the comparison between the bounds related to κ_{CD} will be similar and thus omitted. Since we assume $L_1 = \cdots = L_n$, we have $L \leq \sum_i L_i = nL_1$, and $1 \leq \tau = L/L_1 \leq n$. Therefore, in most cases (more precisely, as long as $\tau \in [1, n/\log^2 n]$ while the full range of τ is $[1, n]$) the bound B_2 is better than B_1. However, B_2 does not dominate B_1 since for our example B_1 is tight while B_2 is $\mathcal{O}(\log^2 n)$ times worse. One natural guess is that maybe the best bound is $B_3 = \mathcal{O}(n^2\kappa\tau)$, which is better than both B_1 and B_2 and also consistent with our example. Unfortunately, B_3 is probably not the right bound since there exists an example such that the $\log^2 n$ factor is unavoidable [40].

Now we discuss the result by Oswald [40]. The paper [40] establishes an upper bound similar to the second bound in 9. Then the paper constructs an example that "matches" the upper bound; more specifically, in the example both κ and τ are $\mathcal{O}(1)$ while the spectral radius of the iteration matrix of C-CD is $1 - 1/\mathcal{O}(\log^2 n)$. Thus the complexity for this example is at least $\mathcal{O}(n^2\log^2 n) = \mathcal{O}(n^2\log^2 n\kappa\tau)$ which is $\log^2 n$ times larger than B_3 [2]. However, this example only "matches" the upper bound in a weak sense as the key parameters κ and τ are constants in the example. In particular, this example has nothing to do with the question whether the extra factor τ is necessary or not. It does not exclude the possibility that the worst-case complexity of C-CD were $\mathcal{O}(n^2\log^2 n\kappa)$ or even $\mathcal{O}(n^2\log^2 n\kappa_{CD})$ which are very close to the complexity of GD and R-CD respectively. We think the extra τ factor is very important for at least two reasons. First, for most randomly generated matrices the ratio $\tau = L/L_1$ is much larger than $\mathcal{O}(\log n)$. This can be tested by numerical experiments, and also validated by theoretical results: for example, for the Wishart random ensamble $A = U^T U$ where the entries of U are standard Gaussian variables, the ratio $\tau = L/L_1$

[2] This statement is not rigorous. It is tricky: the spectral radius of a non-symmetric iteration matrix may not provide the lower bound of the convergence rate; extra effort is needed to rigorously build the connection. We will discuss this issue in more details later.

is approximately $\mathcal{O}(\sqrt{n})$ [3]. Second, τ exactly characterizes the theoretical improvement of R-CD over GD. When τ is small, the gain of using R-CD is very limited: either the problem is too easy and GD already performs well, or the problem is so difficult that even R-CD does not help. Thus the interesting problems for CD-type methods are those with large τ.

Now we know that the $\log^2 n$ factor is necessary for one extreme case $\tau = \mathcal{O}(1)$, and the $\log^2 n$ factor can be removed for the other extreme case $\tau = n$. The transition has to happen somewhere in between, and we guess it happens near $\tau \approx \mathcal{O}(\log^2 n)$. In other words, we guess the "best" bound is

$$B_{\text{conj}} = \mathcal{O}(n^2 \kappa \max\{\tau, C \log^2 n\}).$$

To prove or disprove this bound, we only need to prove or disprove the following conjecture.

Conjecture 3.1. *If A is PSD with all diagonal entries being L_1 and $\tau(A) = \frac{\lambda_{\max}(A)}{L_1} \geq C_1 \log^2 n$ for some constant C_1, then the lower triangular part (with diagonals) Γ, defined as $\Gamma_{ij} = A_{ij}, i \leq j$ and $\Gamma_{ij} = 0, i > j$, satisfies $\|\Gamma\| \leq C_2 \|A\|$.*

In other words, although the operator norm of the triangular operator is $\mathcal{O}(\log n)$, here we conjecture that when restricted to a certain class of PD matrices (τ is not too small) the operator norm becomes $\mathcal{O}(1)$.

3.2 $1/L_{\min}$ factor, Jacobi Preconditioning and Open Questions

We will discuss the complexity bounds when the diagonal entries L_i's are not equal. It turns out the "true" complexity in this general case is much more subtle than the equal-diagonal case (i.e. the case where all L_i's are equal) and related to an old problem in numerical linear algebra.

In the previous discussions we often assume $L_i = L_1, \ \forall i$ since one can always scale the coefficient matrix A to get a new matrix $D_A^{-1/2} A D_A^{-1/2}$ and modify the algorithm correspondingly. Such a preprocessing procedure is called Jacobi preconditioning in numerical linear algebra. It is very simple to implement and only slightly increases the total complexity of the algorithm.

Nevertheless, one may still wonder what the complexity in the non-equal-diagonal case is. Our Proposition (3.1) implies an upper bound which is more general than (9):

$$\min\left\{\tilde{\mathcal{O}}\left(n^3 \kappa \frac{L_{\text{avg}}}{L_{\min}}\right), \tilde{\mathcal{O}}\left(n^2 \kappa \log^2 n \frac{L}{L_{\min}}\right)\right\}. \tag{14}$$

Notice that $L_{\min}$ appears in the denominator of both bounds, thus as $L_{\min} \to 0$ both bounds approach infinity. Intuitively, this implies that when one coordinate has very little contribution to the whole function C-CD will converge very slowly. However, this phenomenon will not happen in practice and the dependency in $1/L_{\min}$ is somewhat artificial. In fact, theoretically we can prove a stronger upper bound of C-CD that does not depend on $1/L_{\min}$, but instead depends on a new condition number.

Proposition 3.2. *(Stronger Upper Bounds) Consider the same setting as Proposition 3.1. We have*

$$f(x^{k+1}) - f^* \leq \min\left\{1 - \frac{1}{n\hat{\kappa}}, 1 - \frac{1}{\hat{L}(2 + \log n/\pi)^2} \frac{1}{\hat{\kappa}}\right\}(f(x^k) - f^*). \tag{15a}$$

$$f(x^{k+1}) - f^* \leq \min\left\{1 - \frac{1}{n^2 \hat{\kappa}_{CD}}, 1 - \frac{1}{\hat{L}^2(2 + \log n/\pi)^2} \frac{1}{\hat{\kappa}_{CD}}\right\}(f(x^k) - f^*). \tag{15b}$$

[3] According to [52, Proposition 6.1] the maximum eigenvalue is about $4n$ and the diagonal entries are the lengths of n-dimensional random vectors which are $\mathcal{O}(\sqrt{n})$, thus after scaling the diagonal entries $\tau = L/L_1 \approx \mathcal{O}(\sqrt{n})$.

Here, the parameters $\hat{L} = \lambda_{\max}(\hat{A})$, $\hat{\kappa} = \lambda_{\max}(\hat{A})/\lambda_{\min}(\hat{A})$, *and* $\hat{\kappa}_{CD} = 1/\lambda_{\min}(\hat{A})$, *where* $\hat{A} = D_A^{-1/2} A D_A^{-1/2}$ *is the Jacobi-preconditioned matrix, and* D_A *is a diagonal matrix consisting of all diagonal entries of* A.

The proof of Proposition 3.2 will be given in Appendix A.2. The proof is almost the same as the proof of Proposition 3.1 except that we should replace matrix A and its lower triangular part Γ by the Jacobi-preconditioned versions. By taking a closer look into the proof, we find that Jacobi-preconditioning is naturally "embedded" in C-CD [4]. This is not surprising since in the update rule (3) we need to scale the diagonals A_{ii} at each step, which is similar to Jacobi-preconditioning (but not the same). Therefore, we can think of $\hat{\kappa}$ as a more appropriate parameter to characterize the complexity of C-CD than the original condition number κ.

Proposition 3.2 implies the following upper bound

$$\min\left\{\tilde{\mathcal{O}}\left(n^3\hat{\kappa}\right), \tilde{\mathcal{O}}\left(n^2\hat{\kappa}\log^2 n\frac{\hat{L}}{\hat{L}_1}\right)\right\}, \tag{16}$$

where $\hat{L}_1 = 1$ since all diagonal entries of matrix $\hat{A}$ are 1. This upper bound reduces to the upper bound (9) for the equal-diagonal case since $\hat{A} = A/L_1$ in that case. Comparing this bound with (14) which also holds for the non-equal-diagonal case, we find that the $L_{\text{avg}}/L_{\min}$ factor disappears here (since this ratio equals 1 for the Jacobi-preconditioned matrix). This can be explained as that the factor $L_{\text{avg}}/L_{\min}$ is absorbed into the new condition number $\hat{\kappa}$; in fact, it is straightforward to prove

$$\hat{\kappa} \le \kappa\frac{L_{\max}}{L_{\min}},$$

thus the upper bound $\tilde{\mathcal{O}}\left(n^3\hat{\kappa}\right)$ immediately implies an upper bound $\tilde{\mathcal{O}}\left(n^3\kappa\frac{L_{\max}}{L_{\min}}\right)$ that is slightly weaker than the first bound in (14). It is not easy to explicitly compare the second bound of (16) and the second bound of (14).

The reason we still present the bound dependent on κ, instead of only presenting the bound dependent on $\hat{\kappa}$, is because the former bound allows us to compare cyclic CD with GD. With the new bound (16), a natural question is how to transfer it to a bound that only depends on the parameters of the original matrix, such as κ. This is related to the following classical question on Jacobi-preconditioning:

What is the relation between the condition number of A and that of the Jacobi-preconditioned matrix?

Intuitively, larger discrepancy in the diagonal entries leads to a larger condition number, thus Jacobi-preconditioning which makes the diagonals equal should reduce the condition number. In other words, one may expect that $\hat{\kappa} \le \kappa$ holds for most of the time, if not always. Unfortunately, it is only known that the relation $\hat{\kappa} \le \kappa$ holds for some special A (more precisely, when A satisfies Young's property (A)[5]), according to Forsythe and Straus [54]. For our purpose, the exact relation $\hat{\kappa} \le \kappa$ is not necessary as we are more interested in the upper bound of $\hat{\kappa}/\kappa$. There are some simple bounds (see, e.g., [55, Lemma 3.2, Lemma 3.3]):

$$\hat{\kappa} \le \kappa \cdot \min\{n, \frac{L_{\max}}{L_{\min}}\}.$$

[4] The Jacobi-preconditioning is also embedded in R-CD, but if we pick the coordinates with probability proportional to A_{ii}, the preconditioning effect disappears.

[5] If the rows and columns of a $(p+q)$-dim matrix can be rearranged so that the upper $p \times p$ and lower $q \times q$ submatrices are diagonal, then the matrix is said to have Property (A) [53]. For example, the tridiagonal matrix satisfies Property (A). Also note that the question in [54] appears in a different form: when is $\hat{A}$ the best conditioned matrix out of all possible diagonally scaled matrix of A?

As a direct corollary, the first bound of (16) implies two upper bounds

$$\min\left\{\tilde{\mathcal{O}}\left(n^4\kappa\right), \tilde{\mathcal{O}}\left(n^3\kappa\frac{L_{\max}}{L_{\min}}\right)\right\}. \tag{17}$$

We have already seen a variant of the above second bound in (14).

What is more interesting is the first bound in (17) $\tilde{\mathcal{O}}\left(n^4\kappa\right)$, which is n^2 times worse than GD, and n times worse than the equal-diagonal case! If we want to express the complexity of cyclic CD purely in terms of κ for the non-equal-diagonal case, $\tilde{\mathcal{O}}\left(n^4\kappa\right)$ is the best upper bound we have right now. There is an $\mathcal{O}(n)$-factor gap between this upper bound and the lower bound $\mathcal{O}(n^3\kappa)$. We believe this gap is artificial and there should be a stronger proof that establishes an upper bound of $\mathcal{O}(n^3\kappa)$. Such a stronger upper bound might be achieved by proving a constant upper bound of $\kappa/\hat{\kappa}$. We pose two open questions:

Open Question 1: Is there a non-equal-diagonal example that cyclic CD has complexity worse than $\tilde{\mathcal{O}}\left(n^3\kappa\right)$? If yes, what about $\tilde{\mathcal{O}}\left(n^4\kappa\right)$?

Open Question 2: Is there a constant upper bound on $\hat{\kappa}/\kappa$, where κ and $\hat{\kappa}$ are the condition numbers of A and the Jacobi-preconditioned $\hat{A}$ respectively? If not, what is the best upper bound of $\hat{\kappa}/\kappa$? Is there an example that the ratio $\hat{\kappa}/\kappa$ achieves $\mathcal{O}(n)$?

We stress again that $\tilde{\mathcal{O}}\left(n^3\hat{\kappa}\right)$ is a tight bound in general, and $\tilde{\mathcal{O}}\left(n^3\kappa\right)$ is a tight bound when A has equal diagonal entries. Thus the above Question 1 is only valid when we consider non-equal-diagonal matrix A and insist on expressing the complexity in terms of the condition number of the original matrix. In some sense, it is not as essential as the question whether there is an $\mathcal{O}(n^2)$ gap between C-CD and R-CD studied in this paper. Nevertheless, it is still a valid question, and becomes more interesting due to its relation to Jacobi-preconditioning.

3.3 Precise Comparison of Time Complexity

In our previous comparison between C-CD and GD/R-CD we have ignored the constants, and we do not state the comparison in a formal result. Next we will formally compare them and quantify the exact gap in terms of the time complexity. In the first result the error is measured in the objective values. In the second result the error is measured in iterates, which allows us to add RP-CD into the comparison and get a better bound for R-CD.

The first proposition shows that to achieve any given relative error in objective values, C-CD takes at least $n/20$ times more operations than GD, and $n^2/40$ times more operations than R-CD. The proof of Proposition 3.3 will be given in Appendix C.1.

Proposition 3.3. *(Compare C-CD with GD, R-CD; objective error) Let $k_{\mathrm{CCD}}(\epsilon)$, $k_{\mathrm{GD}}(\epsilon)$ and $k_{\mathrm{RCD}}(\epsilon)$ be the minimum number of iterations [6] for C-CD, GD , R-CD and RP-CD to achieve (expected) relative error*

$$\frac{E(f(x^k)-f^*)}{f(x^0)-f^*} \leq \epsilon$$

for all initial points in $\mathbb{R}^n$ (for C-CD and GD the expectation operator can be ignored). There exists a

[6] For a fair comparison, here one iteration of CD or RP-CD means one cycle of all coordinates, and for R-CD one iteration means randomly selecting coordinates for n times.

quadratic problem such that

$$\frac{k_{\mathrm{CD}}(\epsilon)}{k_{\mathrm{GD}}(\epsilon)} \geq \frac{n}{2\pi^2} \approx \frac{n}{20}, \tag{18a}$$

$$\frac{k_{\mathrm{CCD}}(\epsilon)}{k_{\mathrm{RCD}}(\epsilon)} \geq \frac{n^2}{4\pi^2} \approx \frac{n^2}{40}. \tag{18b}$$

Remark: It seems that the comparison of C-CD and R-CD is not fair since for R-CD we record the *expected* number of iterations. Nevertheless, it is easy to prove that to guarantee the same error with probability $1-\delta$, we only need $\log(1/\delta)$ times more iterations. For simplicity, we just consider the expected number of iterations of R-CD.

In the above Proposition 3.3, the relative error is defined for the function values; next, we prove a result in which the relative error is defined for the (expected) iterates. The proof of Proposition 3.4 is given in Appendix C.2.

Proposition 3.4. *(Compare C-CD with GD,R-CD and RP-CD; iterates error) Let $K_{\mathrm{CCD}}(\epsilon)$, $K_{\mathrm{GD}}(\epsilon)$, $K_{\mathrm{RCD}}(\epsilon)$ and $K_{\mathrm{RPCD}}(\epsilon)$ be the minimum number of iterations [7] for C-CD, GD , R-CD and RP-CD to achieve (expected) relative error*

$$\frac{\|E(x^k)-x^*\|^2}{\|x^0-x^*\|^2} \leq \epsilon$$

for all initial points in $\mathbb{R}^n$ (for C-CD and GD the expectation operator can be igonred). There exists a quadratic problem such that

$$\frac{K_{\mathrm{CCD}}(\epsilon)}{K_{\mathrm{GD}}(\epsilon)} \geq \frac{n}{2\pi^2} \approx \frac{n}{20}, \tag{19a}$$

$$\frac{K_{\mathrm{CCD}}(\epsilon)}{K_{\mathrm{RCD}}(\epsilon)} \geq \frac{n^2}{2\pi^2} \approx \frac{n^2}{20}, \tag{19b}$$

$$\frac{K_{\mathrm{CCD}}(\epsilon)}{K_{\mathrm{RPCD}}(\epsilon)} \geq \frac{n(n+1)}{2\pi^2} \approx \frac{n(n+2)}{20}. \tag{19c}$$

We present the result for two reasons. First, the convergence of iterates is of interest in some scenarios. Second, we can obtain stronger bounds. In particular, the ratio we obtained for the squared iterates of R-CD is twice as large as that for the function values of R-CD ($n^2/20$ v.s. $n^2/40$). Moreover, we are able to add RP-CD into comparison for the iterates error. We do not include RP-CD in Proposition 3.4 since it seems difficult to compute the convergence rate of the objective error for RP-CD. Despite the advantages, we need to emphasize that the convergence of expected iterate error is a weaker notion of convergence than the convergence of objective error, because the former does not lead to a high probability convergence rate while the latter does (which is because $f(x^k)-f^* \geq 0$). If we could bound $E(\|x^k-x^*\|^2)$ instead of $\|E(x^k-x^*)\|^2$, then high probability convergence rate could also be automatically established; but we are unable to bound $E(\|x^k-x^*\|^2)$ for RP-CD either.

Our theory shows that there exists one example $A=A_c$ such that C-CD takes at least $\frac{n}{2\pi^2} \approx \frac{n}{20}$ times more iterations than GD and $\frac{n^2}{2\pi^2} \approx \frac{n^2}{20}$ times more iterations than R-CD to achieve any accuracy ϵ. While the theory is only established for the case c is very close to 1 (recall c is the off-diagonal entry), we will show in simulations that the predicted gaps do really exist for a wide range of c. Note that "the number of required iterations" is defined for "all initial points" (in other words, "worst-case" initial points). We will show in simulations that even for random initialization the gaps observed in practice match those predicted by Proposition 3.4.

[7] Again, for a fair comparison, here one iteration of CD or RP-CD means one cycle of all coordinates, and for R-CD one iteration means randomly selecting coordinates for n times.

4 Proof of Theorem 3.1

4.1 Difficulties

In general, to prove a lower complexity bound, one only needs to construct an example and compute the convergence rate of the example. However, in our case, computing the convergence rate of the example is not easy due to (at least) two reasons.

First, Gauss-Seidel method can be written as a matrix recursion and its convergence rate is related to the spectral radius of the update matrix. It turns out that the spectral radius of our example does not have a closed form expression; in fact, the spectral radius depends on the roots of an n-th order equation. To resolve this issue, we notice that as the constructed matrix tends to singular (i.e. the off-diagonal entries tend to 1) the n-th order equation will become simple; based on this fact, we are able to bound the spectral radius asymptotically (as off-diagonal entries tend to 1, but still for fixed n).

Second, the update matrix of Gauss-Seidel method is a non-symmetric matrix. A simple, though usually ignored, fact is that for non-symmetric matrix recursion, the spectral radius of the iteration matrix is not the lower bound of the convergence factor in the *real* domain. Note that if we were allowed to pick initial points in the complex domain, then the spectral radius did provide a lower bound of the convergence rate; but here we are only interested in the real initial points. We have not seen a general method to deal with this issue; fortunately, the example we constructed happens to exhibit some special structure so that we can provide a lower bound of the convergence rate. We will discuss this difficulty in more details in Section 4.1.1. There is actually one more difficulty caused by the non-symmetry of the iteration matrix: it is even harder to bound the function error. Fortunately again, we are able to resolve this difficulty due to another special property of the problem. See more details in Step 3 of the outline in Section 4.2.

There is a subtle difference between the lower bound and the upper bound, resulting in the issue of non-symmetry and thus a more involved proof for the lower bound, even if it is just for a specific example. The issue of non-symmetry does not appear in the proof of the upper bound in Proposition C.1 because a symmetrization technique is used. Assuming $x^* = 0$, we need to compute the convergence rate of $f(x^k) = (x^k)^T A(x^k) = \|y^k\|^2$, where $y^k = Ux^k$ in which U satisfies $A = U^T U$. It is easy to get the matrix recursion $y^{k+1} = (I - U\Gamma^{-1}U)y^k$, thus one needs to bound the spectral radius of $M_f = I - U\Gamma^{-1}U$. The spectral radius of a non-symmetric matrix is not easy to directly bound, thus in that proof we instead upper bound the spectral norm $\|M_f\| = \sqrt{\|M_f^T M_f\|}$, which gives a upper bound of $\rho(M_f)$ However, the relaxation from $\rho(M_f)$ to $\|M_f\|$ is not reversible; in other words, even if we prove that for our example $\|M_f\|$ is large, this does not mean $\rho(M_f)$ is large (or C-CD is slow). Thus we have to consider the original non-symmetric form $U\Gamma^{-1}U$ or $\Gamma^{-1}A$ for the lower bound.

4.1.1 Why Non-symmetric Iteration Matrix Causes Difficulty

We discuss why the spectral radius of a non-symmetric iteration matrix does not necessarily lead to a lower bound of the convergence rate (for real initial points). Consider the following matrix recursion

$$y^{k+1} = My^k. \tag{20}$$

We say a sequence $\{y^k\}$ converges with convergence rate τ if $\|y^k\| \leq C\tau^k$, where C is a constant.

A basic result is that if M is symmetric the convergence rate of $\|y^k\|$ is exactly $\rho(M)$. How to prove this result? For the lower bound (i.e. the convergence rate is at least $\rho(M)$), we need to pick the initial point to be the eigenvector of M corresponding to $\rho(M)$. This proof no longer works for non-symmetric

M since its eigenvectors may be complex vectors. One way to resolve this issue is to pick the real part of the complex eigenvector; however, this approach requires additional assumptions to work. More specifically, suppose $Mv = \lambda v$, where $\lambda = \rho(M) = |\lambda|e^{i\theta}$, and pick the initial point $y^0 = \mathrm{Re}(v) = \frac{1}{2}(v + \bar{v})$. The update (20) leads to

$$y^k = M^k y^0 = \frac{1}{2}M^k(v + \bar{v}) = \frac{1}{2}(\lambda^k v + \bar{\lambda}^k \bar{v}) = \mathrm{Re}(\lambda^k v) = |\lambda|^k \mathrm{Re}(e^{ik\theta} v).$$

Suppose $v = (r_1 e^{i\phi_1}, \ldots, r_n e^{i\phi_n})$, then

$$\|y^k\| = \rho(M)^k \sqrt{r_1^2 \cos^2(k\theta + \phi_1) + \cdots + r_n^2 \cos^2(k\theta + \phi_n)}. \tag{21}$$

For the lower bound, we want to prove

$$\|y^k\| \geq C\rho(M)^k, \ \forall k, \tag{22}$$

where C is a constant. Without any additional assumption, this is impossible: if $\phi_j = 0, r_j = 1, \forall j$ and $k\theta = \frac{\pi}{2} + 2m\pi$ for some integer m, then $\|y^k\| = 0$. Intuitively, when all ϕ_j's are close to each other, it is hard to lower bound $\|y^k\|$; but if all ϕ_j's are evenly spread out, then $\|y^k\|$ can be lower bounded. For our problem, it turns out the phase ϕ_j goes to $2j\pi/n$ as c goes to 1, which is the the nicest case we can expect (the phases are equally spaced). In such a nice case, we are able to give a simple lower bound of $\|y^k\|$.

One might wonder whether it is easy to obtain a lower bound in the general case under mild assumptions. We consider the simplest case $n = 2$. If $r_1 = r_2 > 0$ and $0 < |\phi_1 - \phi_2| < \pi/2$, then $\|y^k\| \geq \rho(M)^k r_1 |\sin((\phi_1 - \phi_2)/2)|$. However, if $r_1 \neq r_2$, then even for $n = 2$ we need more assumptions to find a lower bound. Such assumptions can be a relation between r_1/r_2 and $\phi_1 - \phi_2$, which look non-intuitive and seem to be constructed merely for theory. Moreover, it is hard to express the corresponding bound (e.g. $r_1|\sin((\phi_1 - \phi_2)/2)|$) as a function of simple parameters of the original problem. From a practical point of view, the "constant" $r_1|\sin((\phi_1 - \phi_2)/2)|$ can be so small that it already meets the practical need. These issues will become even more complicated when $n > 2$. As a conclusion, when the iteration matrix is non-symmetric, it seems difficult to lower bound the convergence rate in general.

4.2 Proof Outline

The detailed proof is divided into three steps. We will construct an example $\min_x x^T A_c x$ where the coefficient matrix has diagonal entries 1 and off-diagonal entries $c \in (0, 1)$. Obviously $x^* = 0$ is the unique minimum and $f^* = 0$.

In Step 1, we compute the spectral radius of the iteration matrix asymptotically. More specifically, we show that the eigenvalues of the iteration matrix are given by $\lambda_j = 1 - q_j^n$, where q_j's are the roots of the equation $q^n(1 - c + q) = 1$. While the closed form expression of λ_j is difficult to compute, we observe that as $c \to 1$, q_j's tend to the n-th unit roots. We then prove that as $c \to 1$ the spectral radius of the iteration matrix tends to roughly $1 - \frac{2\pi^2}{n\kappa}$.

In Step 2, we prove that for a certain real initial point x^0, the relative error $\frac{\|x^k - x^*\|^2}{\|x^0 - x^*\|^2}$ is lower bounded by $O\left(\left(1 - \frac{2\pi^2}{n\kappa}\right)^{2k}\right)$. In other words, the sequence $\{\|x^k - x^*\|^2\}$ converges at a rate lower bounded by the spectral radius $1 - \frac{2\pi^2}{n\kappa}$. The initial point we choose is the real part of the eigenvector corresponding to the spectral radius of the iteration matrix $I - L^{-1}A$. A crucial property is that the eigenvector has an expression $(1, q, \ldots, q^n)$ where q is an complex eigenvalue of $L^{-1}A$, thus the phases of the initial elements are roughly $2j/\pi, j = 1, \ldots, n$. This property makes the calculation of the relative error $\frac{\|x^k - x^*\|^2}{\|x^0 - x^*\|^2}$ possible.

In Step 3, we prove that the relative error $\frac{f(x^k)-f^*}{f(x^0)-f^*} = \frac{f(x^k)}{f(x^0)}$ is also lower bounded by $\mathcal{O}\left(\left(1-\frac{2\pi^2}{n\kappa}\right)^{2k}\right)$. Again, the special structure of the example is crucial for this step. Unlike GD method where the iteration matrix $I - \frac{1}{\beta}A$ has the same eigenvectors as A, the iteration matrix of CD method $I - L^{-1}A$ has different eigenvectors from A. As we pick x^0 to be the real part of an eigenvector of $L^{-1}A$, it is not clear a priori how to bound $f(x^k) = (x^k)^T A x^k$ and $f(x^0) = (x^0)^T A x^0$. Of course one can lower bound $f(x^k)$ by $\lambda_{\min}(A)\|x^k\|^2$ and upper bound $f(x^0)$ by $\lambda_{\max}(A)\|x^0\|^2$ to get a lower bound of $\frac{f(x^k)}{f(x^0)}$, but this will introduce an extra factor $\frac{\lambda_{\min}(A)}{\lambda_{\max}(A)} = \frac{1-c}{1-c+cn} \approx \frac{1-c}{n}$ which tends to 0 as $c \to 1$. Thus we need to give a tighter bound of either $f(x^k)$ or $f(x^0)$. We choose to bound $f(x^0)$ differently: it turns out $f(x^0) = (x^0)^T A x^0$ can be upper bounded by $\lambda_{\min}(A)\|x^0\|^2$ plus some negligible term (as $c \to 1$), which makes $\frac{f(x^k)}{f(x^0)}$ very close to $\frac{\|x^k\|^2}{\|x^0\|^2}$. The crucial property here is that for our example, the eigenvector corresponding to the spectral radius of the iteration matrix $I - L^{-1}A$ is very close to the eigenvector corresponding to the minimum eigenvalue of A. Needless to say, this property does not hold for general matrix A.

4.3 Formal Proof of Theorem 3.1

Assume the initial point is up to our choice for now. We will show in the end of the proof how to deal with an arbitrary initial point.

For any constant $c \in (0,1)$, consider minimizing the following quadratic function

$$\min_{x\in\mathbb{R}^n} f(x) \triangleq x^T A_c x, \tag{23}$$

where $A_c \in \mathbb{R}^{n\times n}$ is defined as

$$A_c = \begin{bmatrix} 1 & c & \dots & c \\ c & 1 & \dots & c \\ \vdots & \vdots & \ddots & \vdots \\ c & c & \dots & 1 \end{bmatrix} \tag{24}$$

Simple calculation shows that A_c is a positive definite matrix, with one eigenvalue $1-c$ with multiplicity $n-1$ and one eigenvalue $1-c+cn$ with multiplicity 1. Thus the condition number of the matrix is

$$\kappa = \frac{1-c+cn}{1-c}. \tag{25}$$

The optimum of the problem is $x = (0;0;\dots;0)$. Solving this problem is also equivalent to solving a linear system of equations $Ax = 0$.

Step 1: Computing the spectral radius of the iteration matrix, asymptotically. The following lemma shows that the eigenvalues of the matrix A_c are the roots of a polynomial equation. The proof of Lemma 4.1 is given in Appendix B.1.

Lemma 4.1. *Suppose $A = A_c$ is defined by (24) and Γ is the lower triangular part of A_c (with diagonals), and denote $\hat{c} = 1-c$. Suppose the $n+1$ roots of*

$$q^n(q-1+c) = cq \tag{26}$$

are $q_0, q_1, \dots, q_{n-1}, q_n$ among which $q_0 = 0, q_n = 1$, then

$$\lambda_k = \frac{(1-c)(1-q_k)}{1-c-q_k} = 1 - q_k^n,\ k = 0,1,\dots,n-1 \tag{27}$$

are all n eigenvalues of $Z = \Gamma^{-1}A$.

Note that $\lambda_n = 1 - q_n^n = 0$ is not an eigenvalue of Z. Eliminating a factor of q in (26), we have that $q_1, \ldots, q_{n-1}, q_n = 1$ are the n roots of the equation $q^{n-1}(q-1+c) = c$. As c goes to 1, the equation becomes $q^n = 1$, thus the roots q_k will converge to an n-th root of unity. Without loss of generality, we can assume

$$\lim_{c\to 1} q_k = e^{i2k\pi/n}. \tag{28}$$

Obviously $\lim_{c\to 1} \lambda_k = 0, \forall k$.

Next we prove

$$J_k \triangleq \lim_{c\to 1} \frac{1/\kappa}{1-|1-\lambda_k|} = \frac{1}{2n\sin^2(k\pi/n)},\ k = 1, \ldots, n-1. \tag{29}$$

For notational convenience, let $\hat{c} \triangleq 1 - c, \hat{\lambda}_k = 1 - \lambda_k$. Then we have

$$\hat{\lambda}_k \overset{(27)}{=} 1 - \frac{\hat{c}(1-q_k)}{\hat{c}-q_k} = \frac{(\hat{c}-1)q_k}{\hat{c}-q_k} = \frac{cq_k}{q_k-\hat{c}} = \frac{c}{1-\hat{c}/q_k}. \tag{30}$$

Then

$$\begin{aligned} J_k \overset{(25)}{=} \lim_{\hat{c}\to 0} \frac{\hat{c}}{\hat{c}+cn}\frac{1}{1-|\hat{\lambda}_k|} = \frac{1}{n}\lim_{\hat{c}\to 0}\frac{\hat{c}}{1-|\hat{\lambda}_k|} &\overset{(30)}{=} \frac{1}{n}\lim_{\hat{c}\to 0}\frac{\hat{c}}{1-\left|\frac{c}{1-\hat{c}/q_k}\right|} \\ &= \frac{1}{n}\lim_{\hat{c}\to 0}\frac{\hat{c}|1-\hat{c}/q_k|}{|1-\hat{c}/q_k|-c}. \end{aligned} \tag{31}$$

Since $\lim_{\hat{c}\to 0}|1-\hat{c}/q_k| = 1$, from the above relation we have

$$\begin{aligned} J_k &= \frac{1}{n}\lim_{\hat{c}\to 0}\frac{\hat{c}}{|1-\hat{c}/q_k|-c} = \frac{1}{n}\lim_{\hat{c}\to 0}\frac{\hat{c}(|1-\hat{c}/q_k|+c)}{|1-\hat{c}/q_k|^2-c^2} = \frac{2}{n}\lim_{\hat{c}\to 0}\frac{\hat{c}}{|1-\hat{c}/q_k|^2-c^2} \\ &= \frac{2}{n}\lim_{\hat{c}\to 0}\frac{\hat{c}}{1+\hat{c}^2/|q_k|^2-2\hat{c}Re(1/q_k)-c^2} \\ &= \frac{2}{n}\lim_{\hat{c}\to 0}\frac{\hat{c}}{\hat{c}(1+c)+\hat{c}^2/|q_k|^2-2\hat{c}Re(1/q_k)} \\ &= \frac{2}{n}\lim_{\hat{c}\to 0}\frac{1}{1+c+\hat{c}/|q_k|^2-2Re(1/q_k)} \end{aligned} \tag{32}$$

Since $\lim_{\hat{c}\to 0}|q_k| = 1, \lim_{\hat{c}\to 0} Re(1/q_k) = \cos(-2k\pi/n)$, the above relation can be further simplified to

$$J_k = \frac{2}{n}\frac{1}{1+1+0-2\cos(-2k\pi/n)} = \frac{1}{2n\sin^2(k\pi/n)},$$

which proves (29).

Step 2: Bound the relative iterates error.

To simplify the notations, let $q = q_1$ and $\lambda = \lambda_1 = \frac{\hat{c}-\hat{c}q_1}{\hat{c}-q_1}$ from now on.

According to the proof of Lemma 4.1, $\tilde{v} = (\tilde{v}_1; \ldots; \tilde{v}_n)$ is an eigenvector of $Z = \Gamma^{-1}A$ corresponding to λ, where

$$\tilde{v}_j = \frac{c}{\lambda-\hat{c}}q^{j-1}, j = 1, \ldots, n. \tag{33}$$

We scale each entry of $\tilde{v}_j$ by a constant $\frac{\lambda-\hat{c}}{c}q$ to get a new vector $v = (v_1, \ldots, v_j)$, where

$$v_l = q^l,\ l = 1, \ldots, n. \tag{34}$$

Obviously v is also an eigenvector of Z corresponding to λ, i.e. $Zv = \lambda v$.

Now pick the initial point $x^0 = \mathrm{Re}(v)$. Suppose

$$q_1 = q = re^{i\theta}, \tag{35}$$

where $i = \sqrt{-1}$, $r > 0$ and $\theta \in [0, 2\pi)$, then

$$x_j^0 = \mathrm{Re}(e^{\sqrt{-1}j\theta}) = \cos(j\theta),\ j = 1, \ldots, n. \tag{36}$$

Since $x^0 = \frac{1}{2}(v + \bar{v})$, and v and $\bar{v}$ are eigenvectors of $M = I - \Gamma^{-1}A$ with eigenvalues $1 - \lambda$ and $1 - \bar{\lambda}$ respectively, we have

$$x^k = M^k x^0 = \frac{1}{2}M^k(v + \bar{v}) = \frac{1}{2}((1-\lambda)^k v + (1-\bar{\lambda})^k \bar{v}) = \mathrm{Re}((1-\lambda)^k v) \overset{(73)}{=} \mathrm{Re}(q^{kn} v).$$

According to (34), the j-th entry of x^k is

$$x_j^k = \mathrm{Re}(q^{kn+l}) = r^{kn+j}\cos(kn + j)\theta,\ j = 1, \ldots, n.$$

Note that $r = |q| \leq 1$ (otherwise C-CD will diverge, but we know from classical results that C-CD always converges for solving our problem), then we have

$$\begin{aligned} \|x^k\|^2 = \sum_{j=1}^{n} r^{2kn+2j}\cos^2(kn + j)\theta &\geq r^{2kn+2n}\sum_j \cos^2(kn + j)\theta \\ &= r^{(2k+2)n}\frac{1}{2}\left(\sum_j \cos(2kn\theta + 2j\theta) + n\right). \end{aligned} \tag{37}$$

To calculate the sum in the above expression, we will need the following standard equality; for completeness, the proof of this claim is given in Appendix B.2.

Claim 4.1. *For any $z, \phi \in \mathbb{R}$, we have*

$$\sum_{j=1}^{n}\cos(z + j\phi) = \frac{\sin(n\phi/2)\cos(z + (n+1)\phi/2)}{\sin(\phi/2)}. \tag{38}$$

Applying (38) to the expression in (37), we have

$$\|x^k\|^2 \geq \frac{1}{2}r^{(2k+2)n}\left(\frac{\sin(n\theta)\cos(2kn\theta + (n+1)\theta)}{\sin(\theta)} + n\right) \geq \frac{1}{2}r^{(2k+2)n}\left(n - \left|\frac{\sin(n\theta)}{\sin(\theta)}\right|\right). \tag{39}$$

Similar to (37) (but bound r^{2j} from above by 1), we have

$$\begin{aligned} \|x^0\|^2 = \sum_{j=1}^{n} r^{2j}\cos^2(j\theta) &\leq \sum_j \cos^2(j\theta) \\ = \frac{1}{2}\left(\sum_j \cos(2j\theta) + n\right) = \frac{1}{2}\left(\frac{\sin(n\theta)\cos((n+1)\theta)}{\sin(\theta)} + n\right) &\leq \frac{1}{2}\left(n + \left|\frac{\sin(n\theta)}{\sin(\theta)}\right|\right). \end{aligned} \tag{40}$$

Combining the above two relations, we have

$$\frac{\|x^k\|^2}{\|x^0\|^2} \geq r^{2kn+2n}\frac{n - |\sin(n\theta)/\sin(\theta)|}{n + |\sin(n\theta)/\sin(\theta)|} = r^{2kn+2n}\omega_c, \tag{41}$$

where

$$\omega_c \triangleq \frac{n - |\sin(n\theta)/\sin(\theta)|}{n + |\sin(n\theta)/\sin(\theta)|}.$$

According to (28), $q_1 = re^{i\theta}$ converges to $e^{i2\pi/n}$ as $c \to 1$, thus $\theta \to 2\pi/n$ and $|\sin(n\theta)/\sin(\theta)| \to 0$ as $c \to 1$, which further implies $\omega_c \to 1$ as $c \to 1$.

Step 3: Bound the relative objective error.

Suppose $A = U^T U$, and denote $y^k = Ux^k$, then

$$\|y^k\|^2 = (x^k)^T U^T U x^k = (x^k)^T A x^k = f(x^k).$$

Note that the minimum eigenvalue of A is $\hat{c} = 1 - c$, thus

$$\|y^k\|^2 = \|Ux^k\|^2 = (x^k)^T A x^k \geq \hat{c}\|x^k\|^2 \overset{(39)}{\geq} \frac{\hat{c}}{2} r^{(2k+2)n} \left(n - \left| \frac{\sin(n\theta)}{\sin(\theta)} \right| \right). \tag{42}$$

We need to give an upper bound of $\|y^0\|^2$. Denote

$$\gamma_j = \mathrm{Re}(q^j) = \cos(j\theta), \quad \Gamma = \sum_{l=1}^{n} \gamma_l. \tag{43}$$

Then the expression of x^0 given in (36) becomes

$$x^0 = (\gamma_1; \ldots; \gamma_n). \tag{44}$$

Since the j-th row of A is $(c, \ldots, c, 1, c, \ldots, c)$ where 1 is in the j-th position, we can compute the j-th entry of Ax^0 as

$$(Ax^0)_j = c \sum_{l=1}^{n} \gamma_l + (1 - c)\gamma_j \overset{(43)}{=} c\Gamma + \hat{c}\gamma_j. \tag{45}$$

Then we have

$$\|y^0\|^2 = (x^0)^T A x^0 = \sum_j \gamma_j (c\Gamma + \hat{c}\gamma_j) = c\Gamma \sum_j \gamma_j + \hat{c} \sum_j \gamma_j^2 = c\Gamma^2 + \hat{c}\|x^0\|^2. \tag{46}$$

We will show that the second term $\hat{c}\|x^0\|^2$ is the dominant term, which will imply that $\frac{\|y^k\|^2}{\|y^0\|^2} \approx \frac{\hat{c}\|x^k\|^2}{\hat{c}\|x^0\|^2}$. To this end, we need to bound Γ^2. By the definition of Γ in (43), we have

$$\Gamma = \sum_j \gamma_j = \mathrm{Re}(\sum_j q^j) = \mathrm{Re}(q \frac{1 - q^n}{1 - q}),$$

thus

$$\Gamma^2 \leq \left| q \frac{1 - q^n}{1 - q} \right|^2 = r^2 \left| \frac{1 - q^n}{1 - q} \right|^2 = r^2 \left| \frac{\lambda}{1 - q} \right|^2 = r^2 \left| \frac{\hat{c}}{\hat{c} - q} \right|^2 = \hat{c}^2 \left| \frac{q}{\hat{c} - q} \right|^2.$$

Substituting the above relation and (40) into (46), we get

$$\begin{aligned} \|y^0\|^2 &\leq c\hat{c}^2 \left| \frac{q}{\hat{c} - q} \right|^2 + \hat{c} \frac{1}{2} \left(n + \left| \frac{\sin(n\theta)}{\sin(\theta)} \right| \right) \\ &= \hat{c}\alpha_c, \end{aligned} \tag{47}$$

where in the last equality we introduce the definition

$$\alpha_c \triangleq c\hat{c}\left|\frac{q}{\hat{c}-q}\right|^2+\frac{1}{2}\left(n+\left|\frac{\sin(n\theta)}{\sin(\theta)}\right|\right).$$

As $c\to 1$, we have $\hat{c}=1-c\to 0$, $\left|\frac{q}{\hat{c}-q}\right|\to 1$ and $\left|\frac{\sin(n\theta)}{\sin(\theta)}\right|\to 0$, thus

$$\alpha_c \to 0+\frac{n}{2}=\frac{n}{2},\quad \text{as } c\to 1. \tag{48}$$

Combining (42) and (47), we get

$$\frac{\|y^k\|^2}{\|y^0\|^2}\geq \frac{\hat{c}}{2}r^{(2k+2)n}\left(n-\left|\frac{\sin(n\theta)}{\sin(\theta)}\right|\right)\frac{1}{\hat{c}\alpha_c}=\beta_c r^{(2k+2)n}, \tag{49}$$

where in the last equality we introduce the definition

$$\beta_c \triangleq \frac{1}{2\alpha_c}\left(n-\left|\frac{\sin(n\theta)}{\sin(\theta)}\right|\right).$$

According to (48) and the fact that $\left|\frac{\sin(n\theta)}{\sin(\theta)}\right|\to 0$ as $c\to 1$, we have

$$\beta_c \to \frac{1}{n}(n-0)=1,\quad \text{as } c\to 1, \tag{50}$$

which implies that for any $\delta>0$, there exists $c_{\mathrm{u},1}<1$ such that

$$\beta_c > 1-\delta,\ \forall c\in(c_{\mathrm{u},1},1). \tag{51}$$

By the relation between λ_1 and q_1 and the definition of r, we have

$$|1-\lambda_1| \overset{(35)}{=} |q_1|^n \overset{(27)}{=} r^n.$$

According to (29), we have $\lim_{c\to 1}\frac{1/\kappa}{1-r^n}=\lim_{c\to 1}\frac{1/\kappa}{1-|1-\lambda_1|}=\frac{1}{2n\sin^2(\pi/n)}>\frac{n}{2\pi^2}$. Therefore, there exists $c_{\mathrm{u},2}<1$ such that $\frac{1/\kappa}{1-r^n}>\frac{n}{2\pi^2},\ \forall\ c\in(c_{\mathrm{u},2},1)$, i.e.

$$r^n > 1-\frac{2\pi^2}{n\kappa},\ \forall\ c\in(c_{\mathrm{u},2},1). \tag{52}$$

For any $\delta>0$, pick $c\in(\max\{c_{\mathrm{u},1},c_{\mathrm{u},2}\},1)$ and substituting (51) and (52) into (49), we obtain

$$\frac{f(x^k)-f^*}{f(x^0)-f^*}=\frac{\|y^k\|^2}{\|y^0\|^2}\geq(1-\delta)\left(1-\frac{2\pi^2}{n\kappa}\right)^{2k+2}.$$

This proves (11a). The bound (11b) immediately follows since for our example

$$\frac{\kappa}{\kappa_{CD}}=\frac{L}{L_{\mathrm{avg}}}=\frac{1-c+cn}{1}\to n$$

as $c\to 1$.

At last, for an arbitrary initial point x^0 our results still hold since C-CD is invariant with respect to the simultaneous shift of the initial point and the space of variables. More specifically, pick $c\in(0,1)$ such that (11) holds and let v be the eigenvector of A_c given in (34). Consider using C-CD to solve the problem

$$\min_z (z-x^0+v)^T A_c (z-x^0+v)$$

starting from x^0. Applying a linear transformation $z=x-x^0+v$, this algorithm becomes C-CD for solving $\min_x x^T A_c x$ starting from v (the optimal solution x^* and optimal value f^* will change accordingly). Applying the result we have proved for this case, we get the desired result for the case with initial point x^0. **Q.E.D.**

5 Numerical Experiments

In this section, we provide numerical experiments of CD, R-CD, RP-CD and GD for solving quadratic minimization problems. Such experiments have been performed (at least partially) many times before, and different comparison results have been reported: some papers present examples that C-CD performs better than R-CD (e.g., [56]), and others present examples that R-CD/RP-CD performs better (e.g. [24]). Nevertheless, instead of simply stating "sometimes C-CD converges faster, sometimes R-CD converges faster", we will show that the size of off-diagonal entries (relative to diagonal entries) affect the performance of C-CD. We summarize our numerical findings below:

1. Cyclic CD is very slow for solving our example (24), as predicted by our theory, even for random initial points and non-asymptotic c (e.g. $c > 0.5$). In addition, the gap between C-CD and GD/R-CD/RP-CD in our simulation matches the theoretical prediction in Proposition 3.4 very well.

2. The ratio $\tau = L/L_{\min}$ is an important indicator of the performance of C-CD. For randomly generated A with fixed diagonal entries 1 (in which case $L_{\min} = 1$), when A has large spectral norm L, C-CD converges much slower than R-CD/RP-CD; when A has small L, C-CD usually converges as fast as (sometimes faster than) R-CD/RP-CD. In practice, $L/L_{\min}$ is closely related to the ratio of the magnitude of the off-diagonal entries over that of the diagonal entries, thus the size of the off-diagonal entries can be a simple indicator of the performance of C-CD.

3. Similar to many experiments in earlier works, we also find that C-CD converges much faster than GD in all cases we test. This is the opposite to the theory based on worst-case analysis. The bizarre gap between theory and practice has motivated our work, but our work does not fill in this gap; this means that new types of analysis might be needed.

We have explained in Section 4.1.1 why the spectral radius might not be the indicator of the convergence rate when the iteration matrix is non-symmetric. Nevertheless, in most practical scenarios we can still roughly regard the spectral radius as an indicator of the worst-case performance of an method; we just needs to be aware that the truth might be different (though unlikely).

We first present simulation results for our example (24).

Our theoretical results are established for the asymptotic case $c \to 1$. We compute $1 - \rho(M)$ where M is the (expected) iteration matrix of C-CD, R-CD, RP-CD and GD for various values of $c \in (0, 1)$. As analyzed before, when M is symmetric (for R-CD, RP-CD, GD) the value $\rho(M)$ represents the convergence rate. When M is non-symmetric (for C-CD), $\rho(M)$ does not necessarily indicate the convergence rate, but we will regard $\rho(M)$ as a rough lower bound of the worst-case convergence rate of C-CD. What we have proved is as $c \to 1$ $\rho(M)$ is indeed the lower bound of the convergence rate; we have not proved the same holds for $c < 1$ but we believe so [8]. In the last three columns, we divide the values $1 - \rho(M)$ of R-CD, RP-CD and GD by the value of C-CD, and the resulting ratio represents how many times faster compared to C-CD. This is because it roughly takes $\log(1/\epsilon)/(1 - \rho(M))$ iterations to achieve relative error ϵ in the worst-case. In the row "1(theory)", we use the theoretical values $n^2/2\pi^2 \approx n^2/20, n(n+1)/2\pi^2 \approx n(n+1)/20, n/2\pi^2 \approx n/20$ according to Proposition 3.4.

Table 1 clearly shows that for $c = 0.8$ the gap between C-CD and other methods is already large, and rather close to the theory value for $c = 1$. In fact, the gap between GD and C-CD for $c = 0.8$ is around 80%

[8] It is possible to rigorously prove that $\rho(M)$ serves as a lower bound for the converge rate for specific $c < 1$. We can compute the phases of each element of the eigenvector of M corresponding to the spectral radius; as long as the phases are not too clustered (i.e. close to each other), we can give a lower bound of (21). However, even though the simulations clearly show that these phases are not clustered, to rigorously prove such an obvious phenomenon seems to require long and careful computation.

Table 1: Comparison of C-CD, R-CD, RP-CD and GD for our example $A = A_c$

c	$1-\rho(M)$, where M is iteration matrix				Ratio over C-CD		
	C-CD	GD	R-CD	RP-CD	GD	R-CD	RP-CD
n = 20							
0.5	7.6e-1	4.8e-1	4.0e-1	5.2e-1	0.63	0.53	0.68
0.8	1.4e-2	1.2e-1	1.8e-2	2.0e-1	0.85	12.6	14.3
0.99	4.98e-4	5.05e-4	1e-2	1.03e-2	1.01	20.0	20.7
1 (theory)	–	–	–	–	1.01	20.2	21.2
n = 100							
0.5	3.8e-3	9.9e-3	0.39	0.50	2.6	103	132
0.8	6.1e-4	2.5e-3	0.18	0.20	4.08	297	328
0.99	2.0e-5	1.01e-4	0.01	0.01	5.02	494	497
1 (theory)	–	–	–	–	5.07	506	512
n = 1000							
0.5	3.9e-5	9.99e-4	0.39	0.50	25.4	9999	12717
0.8	6.2e-6	2.5e-4	0.18	0.20	40.5	29411	32480
0.99	2.01e-7	1.01e-5	0.01	0.01	50.2	49407	49704
1 (theory)	–	–	–	–	50.7	50600	50760

of the theoretical gap for $c = 1$. When $c = 0.99$, the gap is about 99% of the predicted gap. These findings indicate that the gap between GD and C-CD can be uniformly expressed as c times the theoretical gap for $c = 1$; similarly the gap between R-CD/RP-CD and C-CD can be expressed as c^2 times the theoretical gap for $c = 1$. This phenomenon implies that the lower bound (13) is not only true for $L/L_{\min} = n$, but also for many other values of $L/L_{\min}$ (at least for $L/L_{\min} \geq 0.5n$). Nevertheless, a rigorous validation requires a non-asymptotic analysis for a given c, not for $c \to 1$, which seems not easy.

Table 1 only shows the convergence rate of various methods for the *worst* initial points. Now we present some simulation results for *random* initialization. Figure 1 compares the performance of five methods C-CD, cycCGD-small (cyclic CGD with small stepsize $1/\lambda_{\max}(A)$), RP-CD, R-CD and GD, for minimizing $f(x) = x^T Ax$, where $n = 100, A = A_c$ with $c = 0.8$. The left figure shows the first 100 iterations, and the right figure shows 10^4 iterations. In the right figure, the large gap predicted by theory clearly exists: C-CD is about 4 times slower than GD, and GD is about 80 times slower than R-CD/RP-CD (which means C-CD is about 320 times slower than R-CD/RP-CD, matching Table 1). Figure 1 shows that RP-CD is slightly faster than R-CD, which also matches Table 1.

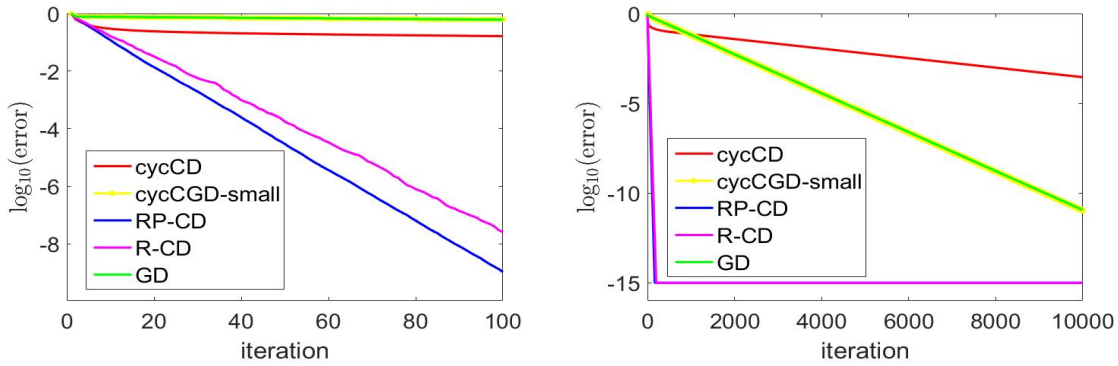

(a) 10^2 iterations (b) 10^4 iterations

Figure 1: Relative error $\frac{f(x^k)-f^*}{f(x^0)-f^*}$ v.s. iterations, for 5 methods C-CD, cyclic CGD with small stepsize $1/\lambda_{\max}(A)$, randomly permuted CD, randomized CD and GD. Minimize $f(x) = x^T Ax$, $n = 100$, $A = A_c$ with $c = 0.8$.

Next, we discuss numerical experiments for randomly generated A; for simplicity, we will normalize the diagonal entries of A to be 1. Since different random distributions of A will lead to different comparison results of various algorithms, we test many distributions and try to understand for which distributions C-CD performs well/poorly. To guarantee that A is positive semidefinite, we generate a random matrix U and let $A = U^T U$. We generate the entries of U i.i.d. from a certain random distribution (which implies that the columns of U are independent), such as $\mathcal{N}(0,1)$ (standard Gaussian distribution), Unif$[0,1]$ (uniform $[0,1]$ distribution), log-normal distribution, etc. It turns out for most distributions C-CD is slower than R-CD and RP-CD, but for standard Gaussian distribution C-CD is better than R-CD and RP-CD.

Inspired by the numerical experiments for the example (24), we suspect that the performance of C-CD depends on how large the off-diagonal entries of A are (with fixed diagonal entries). When $U_{ij} \sim \mathcal{N}(0,1)$, $U^T U$ has small off-diagonal entries since $E(A_{ij}) = 0$, $i \neq j$. When $U_{ij} \sim$ Unif$[0,1]$, $U^T U$ has large off-diagonal entries since $E((U^T U)_{ij}) = n/4$. The difference of these two scenarios lie in the fact that Unif$[0,1]$ has non-zero mean while $\mathcal{N}(0,1)$ has zero mean. It is then interesting to compare the case of drawing entries from Unif$[0,1]$ with the case of drawing entries from Unif$[-0.5, 0.5]$. We will do some sort of A/B testing for each distribution: compare the zero-mean case with the non-zero mean case. To quantify the "off-diagonals over diagonals ratio", we define

$$\chi_i = \frac{\sum_{j\neq i} |A_{ij}|}{A_{ii}},\ i = 1, \ldots, n; \quad \chi_{\text{avg}} = \frac{1}{n}\sum_i \chi_i.$$

We also define

$$\tau = \frac{L}{L_{\min}} = \frac{\lambda_{\max}(A)}{\min_i A_{ii}}.$$

As we have normalized A_{ii} to be 1, we can simply write

$$\chi_i = \sum_{j\neq i} |A_{ij}|, \quad \tau = \lambda_{\max}(A) = L.$$

Obviously $L = \lambda_{\max}(A) \leq 1 + \max_i \chi_i$. In many examples we find L to be close to $1 + \chi_{\text{avg}}$ especially when both of them are large.

We report three pairs of simulation results for three distributions: Gaussian, uniform and log-normal, and each pair consists of a zero-mean case and a non-zero mean case. We summarize the findings from these figures:

1. For all zero-mean cases, C-CD is the fastest; for all non-zero mean cases, C-CD is much slower than R-CD/RP-CD. As mentioned earlier, we suspect the performance of C-CD is closely related to χ_i and χ. With fixed diagonal entries, large off-diagonal entries (or large spectral norm) lead to bad performance of C-CD.

2. Different from the example (24), C-CD is always much faster than GD in these experiments.

3. Overall, RP-CD is the best algorithm out of the five. For the zero-mean case, RP-CD is close to C-CD and about twice faster than R-CD; for the non-zero mean case, RP-CD is $\geq 50\%$ faster than R-CD, and several times faster than GD.

There are many other ways of generating random A. For example, we can always multiply U by the square root of a fixed correlation matrix C. When C has large off-diagonal entries, the results are similar to those shown on the right column of Figure 2. In statistics, this means for solving linear regression problems, C-CD is slow when the data have large correlation ([57] has noticed a related phenomenon). We can also consider non-regression setting, letting U be specially structured matrix such as tri-diagonal matrix or Hankel matrix with entries randomly generated. Similar conclusions hold: either when τ is large we get results similar to the right column of Figure 2, or when τ is small we get the left column.

One interesting question is: why is C-CD much faster than GD in these random experiments, while much slower than GD for our example (24)? Is randomness crucial in the sense that for any random problem C-CD is faster than GD? It turns out the answer is no. We have investigated a related question: is our worst-case example 24 "stable", i.e., if we perturb the example randomly, is C-CD still slow? It is not hard to imagine that C-CD will still be slow if we randomly perturb our example by a very very small amount; the question is how much perturbation can make C-CD faster than GD. It turns out for a reasonable amount of random perturbation C-CD is still slower than GD, and the transition from slow to fast happens rather rapidly. This implies that our "worst-case" example is not like the exponential time example for simplex methods which is very unstable and is later explained by smoothed analysis. Thus for C-CD the smoothed analysis does not help. Knowing that randomization itself is not crucial, the question why C-CD is faster than GD in many cases becomes even more challenging.

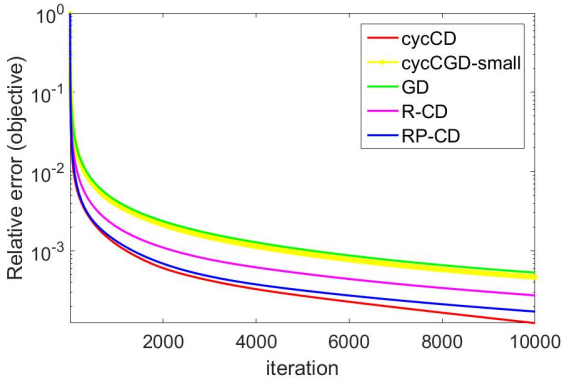

(a1) Gaussian zero mean. $L \approx 3.8$, $\chi_{\max} \approx 7.9$.

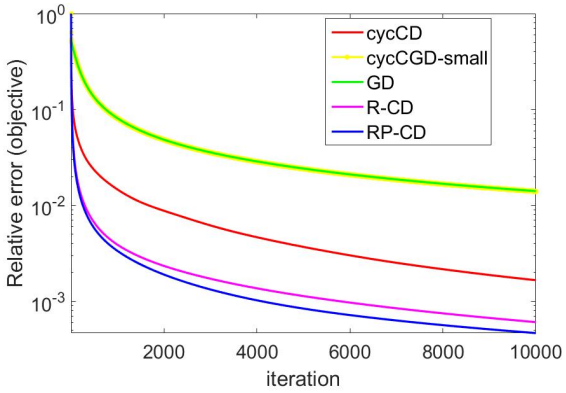

(a2) Gaussian with mean 2. $L \approx 80$, $\chi_{\text{avg}} \approx 79$

.

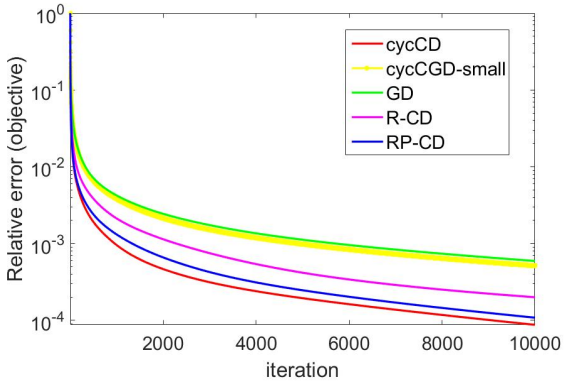

(b1) Uniform [-0.5, 0.5]. $L \approx 3.8$, $\chi_{\text{avg}} \approx 7.9$.

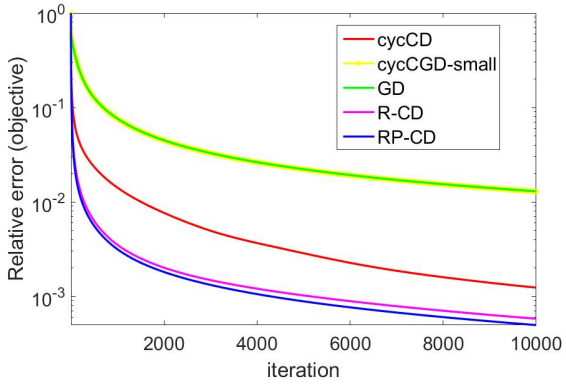

(b2) Uniform [0,1]. $L \approx 75$, $\chi_{\text{avg}} \approx 74$.

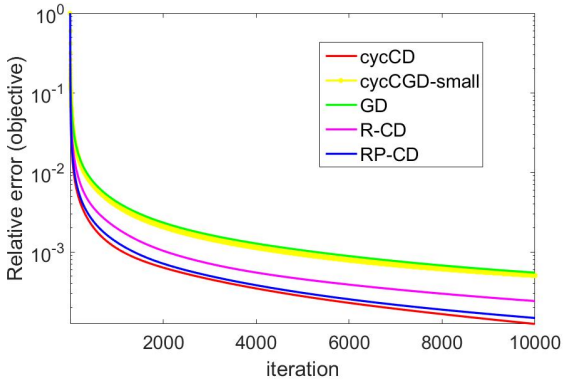

(c1) Log-normal, with zero mean. $L \approx 3.8$, $\chi_{\text{avg}} \approx 7.6$.

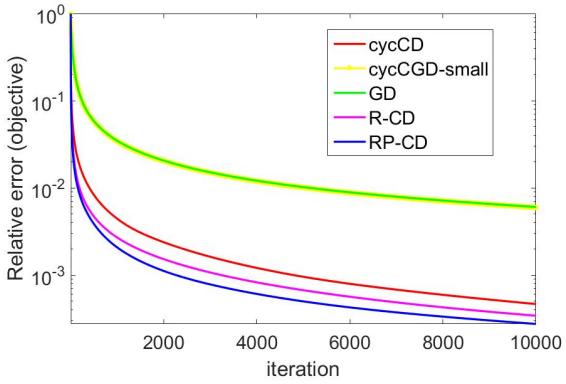

(c2) Log-normal. $L \approx 42$, $\chi_{\text{avg}} \approx 41$.

Figure 2: Comparison of various methods for solving $\min_x x^T Ax$, where A is a 100×100 matrix. A is generated as follows: generate entries of U i.i.d. from a certain distribution and let $A = U^T U$; different figures represent different distributions of U. Figures (a1) and (a2): Gaussian distribution with variance 1; figures (b1) and (b2): uniform distribution; figure (c1) and (c2): log-normal distribution. On the left: zero-mean; on the right: non-zero mean.

Appendix

A Proofs of Upper Bounds

A.1 Proof of Proposition 3.1

Without loss of generality, we can assume $b = 0$. In fact, minimizing $f(x) = x^T Ax - 2b^T x$ is equivalent to minimizing $f(x) = (x - x^*)^T A(x - x^*)$ where $x^* = A^\dagger b$; here we use the fact that $Ax^* = AA^\dagger b = b$ when $b \in \mathcal{R}(A)$. By a linear transformation $z = x - x^*$, C-CD for minimizing $(x - x^*)^T A(x - x^*)$ starting from x^0 is equivalent to C-CD for minimizing $z^T Az$ starting from $z^0 = x^0 - x^*$. Thus we can assume $x^* = 0$, or equivalently, $b = 0$.

The update equation of C-CD now becomes

$$x^{k+1} = (I - \Gamma^{-1}A)x^k = x^k - d^k, \tag{53}$$

where $d^k = \Gamma^{-1}Ax^k$ is the moving direction. This implies

$$\Gamma d^k = Ax^k. \tag{54}$$

We first assume A is positive definite and will show how to extend to the PSD case in the end.

The proof consists of two main claims. The first claim relates the convergence rate of C-CD with the spectral radius of a certain matrix:

Claim A.1. *Let $D_A = diag(A_{11}, \ldots, A_{nn})$ be a diagonal matrix with entries A_{ii}'s. Then*

$$f(x^{k+1}) - f(x^*) \le \left(1 - \frac{1}{\|D_A^{-1/2}\Gamma^T A^{-1}\Gamma D_A^{-1/2}\|}\right)(f(x^k) - f(x^*)). \tag{55}$$

First Proof of Claim A.1 (Optimization Perspective): Following the proof framework of [5], we bound the descent amount and the cost yet to be minimized (cost-to-go) respectively. Suppose $w^0 = x^k, w^n = x^{k+1}$ and $w^1, \ldots, w^{n-1}$ are the $n-1$ intermediate iterates. Since w^i is obtained by minimizing f over the i-th coordinate with other variables fixed, it is easy to verify

$$d_i^k = \frac{1}{2A_{ii}}\nabla_i f(w^{i-1}). \tag{56}$$

In the above expression, $2A_{ii}$ can be viewed as the i-th coordinate-wise Lipschitz constant of ∇f from an optimization perspective. We have

$$w^1 = w^0 - d_1^k e_1 = w^0 - \frac{1}{2A_{11}}\nabla_1 f(w^0), \ldots, w^n = w^{n-1} - d_n^k e_n = w^{n-1} - \frac{1}{2A_{nn}}\nabla_n f(w^{n-1}),$$

where e_i is the i-th unit vector. Then

$$\begin{aligned} f(w^{i-1}) - f(w^i) &= (w^{i-1})^T Aw^{i-1} - (w^{i-1} - d_i^k e_i)^T A(w^{i-1} - d_i^k e_i) = -d_i^k e_i^T Ae_i d_i^k + 2(w^{i-1})^T Ae_i d_i^k \\ &= -A_{ii}(d_i^k)^2 + \nabla_i f(w^{i-1})d_i^k \overset{(56)}{=} -A_{ii}(d_i^k)^2 + 2A_{ii}(d_i^k)^2 = A_{ii}(d_i^k)^2. \end{aligned} \tag{57}$$

Therefore, the descent amount $f(x^k) - f(x^{k+1})$ can be bounded in terms of d^k as

$$f(x^k) - f(x^{k+1}) = \sum_{i=1}^n f(w^{i-1}) - f(w^i) = \sum_{i=1}^n A_{ii}(d_i^k)^2 = (d^k)^T D_A d^k. \tag{58}$$

The cost-to-go estimate is simply

$$f(x^k) - f(x^*) = f(x^k) = (x^k)^T A x^k \overset{(54)}{=} (d^k)^T \Gamma^T A^{-1} \Gamma d^k. \tag{59}$$

Combining with (58), we obtain

$$\frac{f(x^k) - f(x^*)}{f(x^k) - f(x^{k+1})} = \frac{(d^k)^T \Gamma^T A^{-1} \Gamma d^k}{(d^k)^T D_A d^k} \leq \|D_A^{-1/2} \Gamma^T A^{-1} \Gamma D_A^{-1/2}\|, \tag{60}$$

which implies (55).

Second Proof of Claim A.1 (Matrix Recursion Perspective): One natural idea is to prove $f(x^{k+1}) = M_f f(x^k)$ or $f(x^{k+1}) \leq \|M_f\| f(x^k)$ for a certain matrix M_f, based on the update equation of the iterates $x^{k+1} = (I - \Gamma^{-1} A)x^k$. We can write down the expression of $f(x^{k+1})$ in terms of x^k as $f(x^{k+1}) = (x^k)^T (I - \Gamma^{-1} A)^T (I - \Gamma^{-1} A) x^k$. However, it is not clear how this expression is related to $f(x^k) = (x^k)^T A x^k$. A simple trick to resolve this issue is to express everything in terms of d^k. More specifically, we have

$$\begin{aligned} f(x^k) - f(x^{k+1}) &= (x^k)^T A x^k - (x^k - d^k)^T A (x^k - d^k) = 2(d^k)^T A x^k - (d^k)^T A d^k \\ &= 2(d^k)^T \Gamma d^k - (d^k)^T A d^k = (d^k)^T (\Gamma + \Gamma^T) d^k - (d^k)^T A d^k = (d^k)^T D_A d^k, \end{aligned} \tag{61}$$

where the last step is because $\Gamma + \Gamma^T = A + D_A$. Equation (61) is equivalent to (58) derived earlier using another approach. The rest is the same as the first proof. **Q.E.D.**

Remark: Although the second proof seems simpler, for people who are familiar with optimization the first proof is probably easier to understand: equation (57) is just the classical descent lemma (applied to each coordinate), thus (58) is straightforward to derive. In the proof of [5], one crucial step is to bound the cost-to-go in terms of d^k; here for the quadratic case the cost-to-go has a closed form expression given by (60). The second proof is cleaner to write, but it is specifically tailored for the quadratic problem; in contrast, the first proof can be extended to non-quadratic problems as done in [5] ((58) and (60) will become inequalities).

Claim A.2. *Let $D_A = diag(A_{11}, \ldots, A_{nn})$ be a diagonal matrix with entries A_{ii}'s. Then*

$$\|\Gamma^T A^{-1} \Gamma\| \leq \kappa \cdot \min \left\{ \sum_i L_i, (2 + \frac{1}{\pi} \log n)^2 L \right\}. \tag{62}$$

Proof of Claim A.2:

Denote

$$\Gamma_{\text{unit}} = \begin{bmatrix} 1 & 0 & 0 & \ldots & 0 \\ 1 & 1 & 0 & \ldots & 0 \\ 1 & 1 & 1 & \ldots & 0 \\ 1 & 1 & 1 & \ldots & 1 \end{bmatrix},$$

then $\Gamma = \Gamma_{\text{unit}} \circ A$, where $\circ$ denotes the Hadamard product. According to the classical result on the operator norm of the triangular truncation operator [58, Theorem 1], we have

$$\|\Gamma\| = \|\Gamma_{\text{unit}} \circ A\| \leq (1 + \frac{1}{\pi} + \frac{1}{\pi} \log n) \|A\| \leq (2 + \frac{1}{\pi} \log n) \|A\|.$$

Thus we have

$$\|\Gamma^T A^{-1} \Gamma\| \leq \|\Gamma^T \Gamma\| \|A^{-1}\| = \|\Gamma\|^2 \frac{1}{\lambda_{\min}(A)} \leq (2 + \frac{1}{\pi} \log n)^2 \frac{\|A\|^2}{\lambda_{\min}(A)} = (2 + \frac{1}{\pi} \log n)^2 \kappa L,$$

which proves the second part of (62).

We can bound $\|\Gamma\|^2$ in another way (denote λ_i's as the eigenvalues of A):

$$\begin{aligned}\|\Gamma\|^2 \leq \|\Gamma\|_F^2 = \frac{1}{2}(\|A\|_F^2 + \sum_i A_{ii}^2) = \frac{1}{2}\left(\sum_i \lambda_i^2 + \sum_i A_{ii}^2\right) \\ \leq \frac{1}{2}\left((\sum_i \lambda_i)\lambda_{\max} + L_{\max}\sum_i A_{ii}\right) \overset{\text{(i)}}{=} \frac{1}{2}(L + L_{\max})\sum_i L_i \leq L\sum_i L_i.\end{aligned} \tag{63}$$

where(i) is because $\sum_i \lambda_i = \mathrm{tr}(A) = \sum_i A_{ii}$ and $A_{ii} = L_i$. Thus

$$\|\Gamma^T A^{-1}\Gamma\| \leq \|\Gamma\|^2 \frac{1}{\lambda_{\min}(A)} \overset{(63)}{\leq} \frac{L}{\lambda_{\min}}\sum_i L_i = \kappa\sum_i L_i.$$

which proves the first part of (62). **Q.E.D.**

Finally, according to the fact $\|D_A^{-1/2}BD_A^{-1/2}\| \leq \frac{1}{\min_i L_i}\|B\| = \frac{1}{L_{\min}}\|B\|$ for any positive definite matrix B, we have

$$\|D_A^{-1/2}\Gamma^T A^{-1}\Gamma D_A^{-1/2}\| \leq \frac{1}{L_{\min}}\|\Gamma^T A^{-1}\Gamma\| \overset{(62)}{\leq} \frac{1}{L_{\min}}\kappa \cdot \min\left\{\sum_i L_i, (2 + \frac{1}{\pi}\log n)^2 L\right\}$$

Plugging this inequality into (55) and replacing $\sum_i L_i$ by nL_{avg}, we obtain (8a).

Now we show how to modify the above proof to the case that A is PSD. The intuition is that the update equation (53) implies that all $\{x^k\}_{k\geq 1}$ lie in the space $\mathcal{R}(A)$, thus we can view A as a PD matrix when restricted to $\mathcal{R}(A)$. More specifically, from (54) we have

$$x^k = A^\dagger \Gamma d^k.$$

Then (59) is slightly modified to $(x^k)^T A x^k = (d^k)^T\Gamma^T A^\dagger \Gamma d^k$. We still have (58) since its proof does not require A to be positive definite. Now we modify (60) to

$$\begin{aligned}\frac{f(x^k) - f(x^*)}{f(x^k) - f(x^{k+1})} = \frac{(d^k)^T\Gamma^T A^\dagger \Gamma d^k}{(d^k)^T D_A d^k} \leq \frac{(d^k)^T\Gamma^T\Gamma d^k \|A^\dagger\|}{(d^k)^T D_A d^k} \\ \overset{\text{(i)}}{=} \frac{1}{\lambda_{\min}}\frac{(d^k)^T\Gamma^T\Gamma d^k}{(d^k)^T D_A d^k} \leq \frac{1}{\lambda_{\min}}\|D_A^{-1/2}\Gamma^T\Gamma D_A^{-1/2}\| \leq \frac{1}{\lambda_{\min}L_{\min}}\|\Gamma^T\Gamma\|.\end{aligned} \tag{64}$$

where (i) is because $\|A^\dagger\| = 1/\lambda_{\min}$ where $\lambda_{\min}$ is the minimum non-zero eigenvalue of A. The rest is almost the same as the proof for the PD case: obtaining the bounds of $\Gamma^T\Gamma$ as in Claim A.2 and plugging them into (64) immediately leads to (8a).

The first bound of result (8b) is a direct corollary of (8a) because $\kappa \leq n\kappa_{CD}$ (which is because $\lambda_{\max}(A) \leq \mathrm{tr}(A) = nL_{\text{avg}}$). The second bound of (8b) is the same as the second bound of (8a) because

$$\frac{\kappa L}{L_{\min}} = \frac{L^2}{\lambda_{\min}L_{\min}} = \frac{L^2}{L_{\text{avg}}L_{\min}}\frac{L_{\text{avg}}}{\lambda_{\min}} = \frac{L^2}{L_{\text{avg}}L_{\min}}\kappa_{CD}.$$

This finishes the proof of Proposition 3.1.

A.2 Proof of Proposition 3.2

This proof is a slight modification of the proof of Proposition 3.1.

We first consider the case that A is positive definite. The insight is to rewrite the relation proved in Claim A.1

$$f(x^{k+1}) - f(x^*) \leq \left(1 - \frac{1}{\|D_A^{-1/2}\Gamma^T A^{-1}\Gamma D_A^{-1/2}\|}\right)(f(x^k) - f(x^*)) \tag{65}$$

as

$$f(x^{k+1}) - f(x^*) \leq \left(1 - \frac{1}{\|\hat{\Gamma}^T \hat{A}^{-1}\hat{\Gamma}\|}\right)(f(x^k) - f(x^*)), \tag{66}$$

where $\hat{\Gamma} = D_A^{-1/2}\Gamma D_A^{-1/2}$ and $\hat{A} = D_A^{-1/2} A D_A^{-1/2}$. Note that $\hat{\Gamma}$ is still the lower-triangular part (with diagonal entries) of the Jacobi-preconditioned matrix $\hat{A}$. The diagonal entries of $\hat{\Gamma}$ and $\hat{A}$ are all 1, so $\hat{L}_i = 1,\ \forall i$.

Applying Claim A.2 we have

$$\|\hat{\Gamma}^T \hat{A}^{-1}\hat{\Gamma}\| \leq \hat{\kappa} \cdot \min\left\{\sum_i \hat{L}_i, (2 + \frac{1}{\pi}\log n)^2 \hat{L}\right\} = \hat{\kappa} \cdot \min\left\{n, (2 + \frac{1}{\pi}\log n)^2 \hat{L}\right\}.$$

Plugging the above relation into (66) we obtain (15a). Similar to Proposition 3.1, the second bound (15b) follows directly from (15a).

The case that A is PSD is can be handled in a similar way to the proof of Proposition 3.1.

B Supplemental Proofs for Theorem 3.1

B.1 Proof of Lemma 4.1

Suppose λ is an eigenvalue of $Z = L^{-1}A$ and $v = (v_1; v_2; \ldots; v_n) \in \mathbb{C}^{n\times 1}$ is the corresponding eigenvector. Then we have

$$\begin{aligned}
& L^{-1}Av = \lambda v \\
\Longrightarrow\quad & Av = \lambda Lv \\
\Longrightarrow\quad & \begin{cases}
v_1 + c\sum_{j\neq 1} v_j = \lambda v_1 \\
v_2 + c\sum_{j\neq 2} v_j = \lambda(cv_1 + v_2) \\
\cdots \\
v_k + c\sum_{j\neq k} v_j = \lambda(cv_1 + \cdots + cv_{k-1} + v_k), \\
\cdots \\
v_n + c\sum_{j\neq n} v_j = \lambda(cv_1 + \cdots + cv_{n-1} + v_n).
\end{cases}
\end{aligned} \tag{67}$$

Without loss of generality, we can assume

$$\sum_{j=1}^{n} v_j = 1. \tag{68}$$

Let $\hat{c} = 1 - c$. Then (67) becomes

$$\begin{cases} \hat{c}v_1 + c = \lambda v_1 \\ \hat{c}v_2 + c = \lambda(cv_1 + v_2) \\ \dots \\ \hat{c}v_k + c = \lambda(cv_1 + \cdots + cv_{k-1} + v_k), \\ \dots \\ \hat{c}v_n + c = \lambda(cv_1 + \cdots + cv_{n-1} + v_n). \end{cases} \tag{69}$$

The first equation implies $v_1 = \frac{c}{\lambda - \hat{c}}$. Plugging into the second equation, we get

$$v_2 = \frac{c(1 - \lambda v_1)}{\lambda - \hat{c}} = \frac{c(\lambda - \hat{c} - \lambda c)}{(\lambda - \hat{c})^2} = \frac{c\hat{c}(\lambda - 1)}{(\lambda - \hat{c})^2}.$$

Plugging the expression of v_1, v_2 into the third equation, we get

$$v_3 = \frac{c(1 - \lambda v_1 - \lambda v_2)}{\lambda - \hat{c}} = \frac{c(\hat{c})^2(\lambda - 1)^2}{(\lambda - \hat{c})^3}.$$

In general, we can prove by induction that

$$v_k = \frac{c(\hat{c})^{k-1}(\lambda - 1)^{k-1}}{(\lambda - \hat{c})^k} = \frac{c}{\lambda - \hat{c}} q^{k-1}, \tag{70}$$

where

$$q = \frac{\hat{c}(\lambda - 1)}{\lambda - \hat{c}}. \tag{71}$$

We can also express λ in terms of q as

$$\lambda = \frac{\hat{c} - \hat{c}q}{\hat{c} - q}. \tag{72}$$

Note that the expression of v_k given by (70) satisfies (69) for any λ, but our goal is to compute λ. To do this, we need to utilize the normalization assmption (68). In particular, we have (when $q \neq 1$)

$$\begin{aligned} & 1 = \sum_k v_k = (\sum_{k=1}^{n} q^{k-1}) \frac{c}{\lambda - \hat{c}} = \frac{1 - q^n}{1 - q} \frac{c}{\lambda - \hat{c}} \\ \Longrightarrow \quad & (1 - q)(\lambda - \hat{c}) = c(1 - q^n) \\ \overset{(71)}{\Longrightarrow} \quad & c\lambda = c(1 - q^n) \\ \Longrightarrow \quad & q^n = 1 - \lambda \overset{(72)}{=} 1 - \frac{\hat{c} - \hat{c}q}{\hat{c} - q} \\ \Longrightarrow \quad & q^n = \frac{qc}{q - \hat{c}} \\ \Longrightarrow \quad & q^n(q - \hat{c}) = cq. \end{aligned} \tag{73}$$

The above procedure is reversible, i.e. suppose $q \neq 1$ is a root of $q^n(q - \hat{c}) = cq$, then $\lambda = \frac{\hat{c} - \hat{c}q}{\hat{c} - q}$ is an eigenvalue of Z. Suppose the $n + 1$ roots of $q^n(q - \hat{c}) = cq$ are $q_0 = 0, q_1, \dots, q_{n-1}, q_n = 1$ ($q = 0$ and $q = 1$ are always roots), then $\lambda_k = \frac{\hat{c} - \hat{c}q_k}{\hat{c} - q_k} \overset{(73)}{=} 1 - q_k^n, k = 0, \dots, n - 1$ are the n eigenvalues of Z.

B.2 Proof of Claim 4.1

Since $2\sin(n\phi/2)\cos(x+(n+1)\phi/2) = \sin(z+(n+1/2)\phi) - \sin(z+\phi/2)$, the desired equation (38) is equivalent to

$$\sum_{j=1}^{n}\cos(z+j\phi) = \frac{\sin(z+(n+1/2)\phi)-\sin(z+\phi/2)}{2\sin(\phi/2)}. \tag{74}$$

We prove (74) by induction. When $n=1$, it holds because $\sin(z+1.5\phi)-\sin(z+0.5\phi) = 2\sin(\phi/2)\cos(z+\phi)$. Suppose (74) holds for $n-1$, i.e.

$$\sum_{j=1}^{n-1}\cos(z+j\phi) = \frac{\sin(z+(n-1/2)\phi)-\sin(z+\phi/2)}{2\sin(\phi/2)}.$$

Note that $2\cos(z+n\phi)\sin(\phi/2) = \sin(z+(n+1/2\phi)) - \sin(z+(n-1/2)\phi)$, therefore

$$\begin{aligned}\sum_{j=1}^{n}\cos(z+j\phi) &= \frac{\sin(z+(n-1/2)\phi)-\sin(z+\phi/2)}{2\sin(\phi/2)} + \cos(z+n\phi)\\ &= \frac{\sin(z+(n-1/2)\phi)-\sin(z+\phi/2)+\sin(z+(n+1/2\phi))-\sin(z+(n-1/2)\phi)}{2\sin(\phi/2)} \\ &= \frac{\sin(z+(n+1/2)\phi)-\sin(z+\phi/2)}{2\sin(\phi/2)}.\end{aligned} \tag{75}$$

This completes the induction step, and thus (74) holds. **Q.E.D.**

C Proofs of Propositions on Exact Comparison

C.1 Proof of Proposition 3.3

Same as the proof of Theorem 3.1, we pick $A = A_c$ and consider minimizing $f(x) = x^T A_c x$. Obviously the minimizer $x^* = 0$ and the optimal value $f^* = f(x^*) = 0$.

We first compute $k_{\mathrm{GD}}(\epsilon)$. The update equation of GD is $x^k = (I - \frac{1}{\beta}A)x^{k-1}$, where $\beta = \lambda_{\max}(A)$. Since A is a symmetric positive definite matrix, we can assume $A = U^T U$, where $U \in \mathbb{R}^{n\times n}$ is non-singular. Then

$$Ux^k = U(I-\frac{1}{\beta}A)x^{k-1} = (I - \frac{1}{\beta}UU^T)Ux^{k-1}, \tag{76}$$

The spectral norm of the iteration matrix

$$\|I - \frac{1}{\beta}UU^T\| = \|I - \frac{1}{\beta}A\| = 1 - \frac{1}{\beta}\lambda_{\min}(A) = 1 - 1/\kappa, \tag{77}$$

where κ is the condition number of A given by

$$\kappa = \frac{1-c+cn}{1-c}. \tag{78}$$

The relation (76) implies

$$f(x^k) = \|Ux^k\|^2 \le \|I - \frac{1}{\beta}UU^T\|^2\|Ux^{k-1}\|^2.$$

Therefore we have

$$f(x^k) \leq \|I - \frac{1}{\beta} UU^T\|^{2k} f(x^0) = \left(1 - \frac{1}{\kappa}\right)^{2k} f(x^0). \tag{79}$$

The minimum number of iterations to achieve $\frac{(x^k)^H A x^k}{(x^0)^H A x^0} \leq \epsilon$ for all initial points $x^0 \in \mathbb{R}^{n\times 1}$ can be upper bounded as

$$k_{\mathrm{GD}}(\epsilon) \leq \left\lceil \frac{1}{2} \frac{\ln \epsilon}{\|I - A/\beta\|} \right\rceil = \frac{1}{2} \frac{\ln \epsilon}{\ln(1 - 1/\kappa)} + 1. \tag{80}$$

Let $y^k = Ux^k$. We will use the same definitions of q, r, λ_i as in the proof of Theorem 3.1. According to (49), to obtain a relative error $\frac{f(x^k)}{f(x^0)} = \frac{\|y^k\|^2}{\|y^0\|^2} \leq \epsilon$, the number of iterations $k = k_{\mathrm{CCD}}(\epsilon)$ should satisfy

$$\epsilon \geq \beta_c r^{(2k+2)n} \overset{(35)}{=} \beta_c |q|^{(2k+2)n} \overset{(27)}{=} \beta_c |1 - \lambda_1|^{2k+2},$$

i.e.

$$k_{\mathrm{CCD}}(\epsilon) \geq \frac{1}{2} \frac{\ln(1/\epsilon) + \ln(\beta_c)}{\ln(1/|1 - \lambda_1|)} - 1. \tag{81}$$

Since $\lim_{c\to 1} \ln(1/|1 - \lambda_1|) = 0$ and by (50) $\lim_{c\to 1} \beta_c = 1$, we have

$$\lim_{c\to 1} \frac{\frac{\ln(1/\epsilon)+\ln(\beta_c)}{\ln(1/|1-\lambda_1|)} - 1}{\frac{\ln(1/\epsilon)}{\ln(1/|1-\lambda_1)|}} = 1.$$

$$\begin{aligned}
\lim_{c\to 1} \frac{k_{\mathrm{CCD}}(\epsilon)}{k_{\mathrm{GD}}(\epsilon)} &\geq \lim_{c\to 1} \left(\frac{1}{2} \frac{\ln(1/\epsilon) + \ln(\beta_c)}{\ln(1/|1 - \lambda_1|)} - 1\right) \left(\frac{1}{2} \frac{\ln \epsilon}{\ln(1 - 1/\kappa)} + 1\right)^{-1} \\
&= \lim_{c\to 1} \frac{1}{2} \frac{\ln(1/\epsilon)}{\ln(1/|1 - \lambda_1|)} \left(\frac{1}{2} \frac{\ln \epsilon}{\ln(1 - 1/\kappa)}\right)^{-1} \\
&= \lim_{c\to 1} \frac{\ln(1 - 1/\kappa)}{\ln|1 - \lambda_1|} = \lim_{c\to 1} \frac{\ln(1 - 1/\kappa)}{-1/\kappa} \cdot \frac{-(1 - |1 - \lambda_1|)}{\ln|1 - \lambda_1|} \cdot \frac{1/\kappa}{1 - |1 - \lambda_1|} \\
&= \lim_{c\to 1} \frac{1/\kappa}{1 - |1 - \lambda_1|} \overset{(29)}{=} \frac{1}{2n \sin^2(\pi/n)} > \frac{n}{2\pi^2}.
\end{aligned} \tag{82}$$

The convergence rate of the objective values for R-CD has been given in [2, Theorem 2] and [1, Theorem 3.6]. We present the convergence rate of both the iterates and the objective values for R-CD, when solving quadratic problems (23). The proof is quite straightforward and omitted here. Note that the proposition implies $\|E(x^k)\|^2$ converges twice as fast as $E(f(x^k))$, which explains why in Proposition 3.4 the gap between C-CD and R-CD is twice as large as that in Proposition 3.3.

Proposition C.1. *Consider solving a quadratic minimization problem* (23) *where A is a positive definite matrix with all diagonal entries being* 1. *Suppose R-CD generates a sequence z^k according to* (5) *and define $x^k = z^{kn}$. Then*

$$\|E(x^k)\|^2 \leq \left(1 - \frac{1}{n}\lambda_{\min}\right)^{2kn} \|x^0\|^2, \tag{83}$$

and

$$E(f(x^k)) \leq (1 - \frac{1}{n}\lambda_{\min})^{kn} f(x^0), \tag{84}$$

where $\lambda_{\min}$ is the minimum eigenvalue of A.

According to Proposition C.1,

$$E(f(x^k)) \leq (1 - 1/\kappa_{\rm CD})^{kn} f(x^0),$$

where $\kappa_{\rm CD} = \frac{\max_i A_{ii}}{\lambda_{\min}(A)}$. To achieve an error $\frac{E(f(x^k))}{f(x^0)} \leq \epsilon$, we only need

$$(1 - 1/\kappa_{\rm CD})^{kn} \leq \epsilon \Longleftrightarrow k \geq \frac{1}{n} \frac{\ln \epsilon}{\ln(1 - 1/\kappa_{\rm CCD})}.$$

Therefore we have

$$k_{\rm RCD}(\epsilon) \leq \frac{1}{n} \frac{\ln \epsilon}{\ln(1 - 1/\kappa_{\rm CCD})} + 1. \tag{85}$$

Combining the above relation with (80), we have

$$\lim_{c \to 1} \frac{k_{\rm GD}(\epsilon)}{k_{\rm RCD}(\epsilon)} \geq \lim_{c \to 1} \frac{n \ln(1 - 1/\kappa_{\rm CCD})}{2 \ln(1 - 1/\kappa)} = \lim_{c \to 1} \frac{n/\kappa_{\rm CCD}}{2/\kappa} = \lim_{c \to 1} \frac{n \lambda_{\max}(A)}{2 \sum_i \beta_i} = \lim_{c \to 1} \frac{n(1 - c + cn)}{2n} = \frac{n}{2}. \tag{86}$$

Combining the above relation with (82), we obtain

$$\lim_{c \to 1} \frac{k_{\rm CCD}(\epsilon)}{k_{\rm RCD}(\epsilon)} > \frac{n}{2\pi^2} \frac{n}{2} = \frac{n^2}{4\pi^2}. \tag{87}$$

According to (82) and (87), there exists c such that (19a) and (19b) hold. **Q.E.D.**

C.2 Proof of Proposition 3.4

Same as the proof of Theorem 3.1, we pick $A = A_c$ and consider minimizing $f(x) = x^T A_c x$. Obviously the minimizer $x^* = 0$ and the optimal value $f^* = f(x^*) = 0$.

First we consider $K_{\rm GD}(\epsilon)$. Since A is a symmetric positive definite matrix, we can assume $A = U^T U$, where $U \in \mathbb{R}^{n \times n}$ is non-singular. The update formula of GD is $x^{k+1} = (I - \frac{1}{\beta} A)$. The iteration matrix $I - \frac{1}{\beta} A$ has the same eigenvalues as $I - \frac{1}{\beta} U U^T$, the iteration matrix of $\{U x^k\}$ (see (76)). Since both $I - \frac{1}{\beta} A$ and $I - \frac{1}{\beta} U U^T$ are symmetric, we have the following relation (which means that for GD the squared iterates and the function values converge at the same speed)

$$k_{GD}(\epsilon) = K_{GD}(\epsilon). \tag{88}$$

We then consider $K_{\rm CCD}(\epsilon)$. Compare 49 with 41 in the proof of Theorem 3.1, we know that the bound we obtained for the function values is the same as the bound for the squared iterates. Similar to (81), we have

$$K_{\rm CCD}(\epsilon) \geq \frac{1}{2} \frac{\ln(1/\epsilon) + \ln(\omega_c)}{\ln(1/|1 - \lambda_1|)} - 1.$$

Similar to (82) in the proof of Proposition 3.3, we have

$$\lim_{c \to 1} \frac{K_{\rm CCD}(\epsilon)}{K_{\rm GD}(\epsilon)} > \frac{n}{2\pi^2}. \tag{89}$$

Next, we consider $K_{\rm RCD}(\epsilon)$. According to Proposition C.1,

$$\|E(x^k)\| \leq (1 - 1/\kappa_{CD})^{kn} \|x^0\|,$$

which implies

$$K_{\mathrm{RCD}}(\epsilon) \leq \frac{1}{2n} \frac{\ln(\epsilon)}{\ln(1-1/\kappa_{\mathrm{CD}})} + 1.$$

Note that the RHS (right-hand side) of the above bound is asymptotically half the RHS of (85). Combining with (86) and (88), we have

$$\lim_{c\to 1} \frac{K_{\mathrm{GD}}(\epsilon)}{K_{\mathrm{RCD}}(\epsilon)} \overset{(88)}{=} \lim_{c\to 1} \frac{k_{\mathrm{GD}}(\epsilon)}{K_{\mathrm{RCD}}(\epsilon)} = 2 \lim_{c\to 1} \frac{k_{\mathrm{GD}}(\epsilon)}{k_{\mathrm{RCD}}(\epsilon)} \overset{(86)}{>} n.$$

Multiplying this inequality with (89), we have

$$\lim_{c\to 1} \frac{K_{\mathrm{CCD}}(\epsilon)}{K_{\mathrm{RCD}}(\epsilon)} > \frac{n^2}{2\pi^2}. \tag{90}$$

Finally, we compute $K_{\mathrm{RPCD}}(\epsilon)$.

Claim C.1. *Consider using RP-CD (randomly permuted coordinate descent) to solve the problem* $\min_{x\in\mathbb{R}^n} x^T A_c x$ *with* A_c *given in* (24)*. Suppose the initial point is* x^0*, then we have*

$$\|E(x^k)\|^2 \leq (1-(1-c)(1-\gamma))^{2k} \|x^0\|^2, \tag{91}$$

where

$$\gamma = \frac{-n + (1-(1-c)^n)/c}{n(n-1)}. \tag{92}$$

The proof of this claim is given in Appendix C.2.1.

By the definition of γ in (92) we have

$$\lim_{c\to 1} \gamma = -1/n. \tag{93}$$

Similar to the proof of Proposition 3.3, from (91) and (79) we have

$$\begin{aligned}\lim_{c\to 1} \frac{k_{\mathrm{GD}}(\epsilon)}{K_{\mathrm{RPCD}}(\epsilon)} = \lim_{c\to 1} \frac{\kappa}{1/[(1-c)(1-\gamma)]} &\overset{(78)}{=} \lim_{c\to 1} \frac{1-c+cn}{1-c}(1-c)(1-\gamma) \\ &\overset{(92)}{=} \lim_{c\to 1}(1-c+cn)(1-\gamma) \overset{(93)}{=} (n)(1+1/n) = n+1.\end{aligned}$$

Multiplying this relation with (89) and use the fact $K_{\mathrm{GD}}(\epsilon) = k_{\mathrm{GD}}(\epsilon)$ we get

$$\lim_{c\to 1} \frac{K_{\mathrm{CCD}}(\epsilon)}{K_{\mathrm{RPCD}}(\epsilon)} > \frac{n(n+1)}{2\pi^2}. \tag{94}$$

According to (89), (90) and (94), there exists c such that all three relations in (19) hold.

C.2.1 Proof of Claim C.1

For simplicity, we denote $L = L_{12\ldots n}$. Since $A = A_c$, we have

$$L = \begin{bmatrix} 1 & 0 & \ldots & 0 \\ c & 1 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ c & c & \ldots & 1 \end{bmatrix}$$

It is easy to get (recall that $\hat{c} \triangleq 1-c$)

$$L^{-1} = \begin{bmatrix} 1 & 0 & 0 & \dots & 0 & 0 \\ -c & 1 & 0 & \dots & 0 & 0 \\ -c\hat{c} & -c & 1 & \dots & 0 & 0 \\ -c\hat{c}^2 & -c\hat{c} & -c & \ddots & \vdots & \vdots \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\ -c\hat{c}^{n-2} & -c\hat{c}^{n-3} & -c\hat{c}^{n-4} & \dots & -c & 1 \end{bmatrix}$$

Since L_σ can be obtained by permuting the rows and columns of L, thus L_σ^{-1} can also be obtained by similar permutations based on L_σ^{-1}. As the matrix A_c has only two distinct values, we know that the expression of $E(L_\sigma^{-1})$ (the expectation is taken over the uniform distribution of permutations of $\{1,2,\dots,n\}$) must have the following form

$$E(L_\sigma^{-1}) = \begin{bmatrix} 1 & \gamma & \dots & \gamma \\ \gamma & 1 & \dots & \gamma \\ \vdots & \vdots & \ddots & \vdots \\ \gamma & \gamma & \dots & 1 \end{bmatrix}, \tag{95}$$

where γ only depends on c. Due to symmetry, γ must be the average of all off-diagonal entries of L^{-1}, i.e.

$$\gamma = \frac{(n-1)c + (n-2)c\hat{c} + \dots + \hat{c}^{n-2}}{n(n-1)} = \frac{-n + (1-\hat{c}^n)/c}{n(n-1)}. \tag{96}$$

By the expressions (95) and (24), we have

$$E(L_\sigma^{-1})A = \begin{bmatrix} \alpha & \tau & \dots & \tau \\ \tau & \alpha & \dots & \tau \\ \vdots & \vdots & \ddots & \vdots \\ \tau & \tau & \dots & \alpha \end{bmatrix},$$

where $\alpha = 1 + (n-1)c\gamma$, $\tau = c + \gamma + (n-2)c\gamma$. The minimum eigenvalue of this matrix is

$$\lambda_{\min}(E(L_\sigma^{-1})A) = \alpha - \tau = 1 + (n-1)c\gamma - c - \gamma - (n-2)c\gamma = (1-c)(1-\gamma). \tag{97}$$

According to (6), we have (note that σ^k is independent of x^k)

$$E(x^k) = E(I - L_\sigma^{-1}A)E(x^{k-1}),$$

where the expectation is taken over the uniform distribution of permutations of $\{1,2,\dots,n\}$. This implies

$$\|E(x^k)\| \le \|I - E(L_\sigma^{-1})A\| \|E(x^{k-1})\|,$$

which further implies

$$\|E(x^k)\|^2 \le \|I - E(L_\sigma^{-1})A\|^{2k} \|x^0\|^2. \tag{98}$$

It is easy to verify that $I - E(L_\sigma^{-1})A$ is a positive semidefinite matrix, thus $\|I - E(L_\sigma^{-1})A\| = 1 - \lambda_{\min}(E(L_\sigma^{-1})A)$. Plugging this and (97) into (98), we obtain the desired inequality (91).

References

[1] Dennis Leventhal and Adrian S Lewis, "Randomized methods for linear constraints: convergence rates and conditioning," *Mathematics of Operations Research*, vol. 35, no. 3, pp. 641–654, 2010.

[2] Y. Nesterov, "Efficiency of coordiate descent methods on huge-scale optimization problems," *SIAM Journal on Optimization*, vol. 22, no. 2, pp. 341–362, 2012.

[3] A. Beck and L. Tetruashvili, "On the convergence of block coordinate descent type methods," *SIAM Journal on Optimization*, vol. 23, no. 4, pp. 2037–2060, 2013.

[4] M. Hong, X. Wang, M. Razaviyayn, and Z.-Q. Luo, "Iteration complexity analysis of block coordinate descent methods," 2013, Preprint, available online arXiv:1310.6957.

[5] Ruoyu Sun and Mingyi Hong, "Improved iteration complexity bounds of cyclic block coordinate descent for convex problems," in *Advances in Neural Information Processing Systems*, 2015, pp. 1306–1314.

[6] Anne Greenbaum, *Iterative methods for solving linear systems*, vol. 17, Siam, 1997.

[7] Adrian A Canutescu and Roland L Dunbrack, "Cyclic coordinate descent: A robotics algorithm for protein loop closure," *Protein science*, vol. 12, no. 5, pp. 963–972, 2003.

[8] Charles A Bouman and Ken Sauer, "A unified approach to statistical tomography using coordinate descent optimization," *Image Processing, IEEE Transactions on*, vol. 5, no. 3, pp. 480–492, 1996.

[9] Cho-Jui Hsieh, Kai-Wei Chang, Chih-Jen Lin, S Sathiya Keerthi, and Sellamanickam Sundararajan, "A dual coordinate descent method for large-scale linear SVM," in *Proceedings of the 25th international conference on machine learning*. ACM, 2008, pp. 408–415.

[10] Tamara G Kolda and Brett W Bader, "Tensor decompositions and applications," *SIAM review*, vol. 51, no. 3, pp. 455–500, 2009.

[11] Jerome Friedman, Trevor Hastie, and Rob Tibshirani, "Regularization paths for generalized linear models via coordinate descent," *Journal of statistical software*, vol. 33, no. 1, pp. 1, 2010.

[12] Joseph K Bradley, Aapo Kyrola, Danny Bickson, and Carlos Guestrin, "Parallel coordinate descent for l1-regularized loss minimization," *arXiv preprint arXiv:1105.5379*, 2011.

[13] Rahul Mazumder, Jerome H Friedman, and Trevor Hastie, "Sparsenet: Coordinate descent with nonconvex penalties," *Journal of the American Statistical Association*, 2012.

[14] Zaiwen Wen, Donald Goldfarb, and Katya Scheinberg, "Block coordinate descent methods for semidefinite programming," in *Handbook on Semidefinite, Conic and Polynomial Optimization*, pp. 533–564. Springer, 2012.

[15] Meisam Razaviyayn, Mingyi Hong, and Zhi-Quan Luo, "A unified convergence analysis of block successive minimization methods for nonsmooth optimization," *SIAM Journal on Optimization*, vol. 23, no. 2, pp. 1126–1153, 2013.

[16] H. Baligh, M. Hong, W.-C. Liao, Z.-Q. Luo, M. Razaviyayn, M. Sanjabi, and R. Sun, "Cross-layer provision of future cellular networks: A WMMSE-based approach," *IEEE Signal Processing Magazine*, vol. 31, no. 6, pp. 56–68, Nov 2014.

[17] Ruoyu Sun, Hadi Baligh, and Zhi-Quan Luo, "Long-term transmit point association for coordinated multipoint transmission by stochastic optimization," in *Signal Processing Advances in Wireless Communications (SPAWC), 2013 IEEE 14th Workshop on*. IEEE, 2013, pp. 330–334.

[18] Mingyi Hong, Ruoyu Sun, Hadi Baligh, and Zhi-Quan Luo, "Joint base station clustering and beamformer design for partial coordinated transmission in heterogeneous networks," *Selected Areas in Communications, IEEE Journal on*, vol. 31, no. 2, pp. 226–240, 2013.

[19] Ruoyu Sun and Zhi-Quan Luo, "Guaranteed matrix completion via nonconvex factorization," in *2015 IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2015, pp. 270–289.

[20] Z.-Q. Luo and P. Tseng, "On the convergence of the coordinate descent method for convex differentiable minimization," *Journal of Optimization Theory and Application*, vol. 72, no. 1, pp. 7–35, 1992.

[21] D. P. Bertsekas, *Nonlinear Programming, 2nd ed*, Athena Scientific, Belmont, MA, 1999.

[22] P. Tseng, "Convergence of a block coordinate descent method for nondifferentiable minimization," *Journal of Optimization Theory and Applications*, vol. 103, no. 9, pp. 475–494, 2001.

[23] L. Grippo and M. Sciandrone, "On the convergence of the block nonlinear Gauss-Seidel method under convex constraints," *Operations Research Letters*, vol. 26, pp. 127–136, 2000.

[24] Shai Shalev-Shwartz and Tong Zhang, "Stochastic dual coordinate ascent methods for regularized loss," *The Journal of Machine Learning Research*, vol. 14, no. 1, pp. 567–599, 2013.

[25] P. Richtárik and M. Takáč, "Iteration complexity of randomized block-coordinate descent methods for minimizing a composite function," *Mathematical Programming*, vol. 144, pp. 1–38, 2014.

[26] Zhaosong Lu and Lin Xiao, "On the complexity analysis of randomized block-coordinate descent methods," *Mathematical Programming*, vol. 152, no. 1-2, pp. 615–642, 2015.

[27] Zheng Qu, Peter Richtárik, and Tong Zhang, "Randomized dual coordinate ascent with arbitrary sampling," *arXiv preprint arXiv:1411.5873*, 2014.

[28] Qihang Lin, Zhaosong Lu, and Lin Xiao, "An accelerated proximal coordinate gradient method and its application to regularized empirical risk minimization," *arXiv preprint arXiv:1407.1296*, 2014.

[29] Yuchen Zhang and Lin Xiao, "Stochastic primal-dual coordinate method for regularized empirical risk minimization," in *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, 2015, pp. 353–361.

[30] Olivier Fercoq and Peter Richtárik, "Accelerated, parallel, and proximal coordinate descent," *SIAM Journal on Optimization*, vol. 25, no. 4, pp. 1997–2023, 2015.

[31] Ji Liu, Stephen J Wright, Christopher Ré, Victor Bittorf, and Srikrishna Sridhar, "An asynchronous parallel stochastic coordinate descent algorithm," *The Journal of Machine Learning Research*, vol. 16, no. 1, pp. 285–322, 2015.

[32] Andrei Patrascu and Ion Necoara, "Efficient random coordinate descent algorithms for large-scale structured nonconvex optimization," *Journal of Global Optimization*, vol. 61, no. 1, pp. 19–46, 2015.

[33] Cho-Jui Hsieh, Hsiang-Fu Yu, and Inderjit S Dhillon, "Passcode: Parallel asynchronous stochastic dual co-ordinate descent," *arXiv preprint arXiv:1504.01365*, 2015.

[34] Stephen J Wright, "Coordinate descent algorithms," *Mathematical Programming*, vol. 151, no. 1, pp. 3–34, 2015.

[35] Yin Tat Lee and Aaron Sidford, "Efficient accelerated coordinate descent methods and faster algorithms for solving linear systems," in *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*. IEEE, 2013, pp. 147–156.

[36] Zeyuan Allen-Zhu and Lorenzo Orecchia, "Nearly-linear time positive lp solver with faster convergence rate," in *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*. ACM, 2015, pp. 229–236.

[37] Manindra Agrawal, Neeraj Kayal, and Nitin Saxena, "PRIMES is in P," *Annals of mathematics*, pp. 781–793, 2004.

[38] A. Beck, "On the convergence of alternating minimization with applications to iteratively reweighted least squares and decomposition schemes," *SIAM Journal on Optimization*, vol. 25, no. 1, pp. 185–209, 2015.

[39] A. Saha and A. Tewari, "On the nonasymptotic convergence of cyclic coordinate descent method," *SIAM Journal on Optimization*, vol. 23, no. 1, pp. 576–601, 2013.

[40] Peteer Oswald, "On the convergence rate of sor: a worst case estimate," *Computing*, vol. 52, no. 3, pp. 245–255, 1994.

[41] Daniel A Spielman and Shang-Hua Teng, "Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time," *Journal of the ACM (JACM)*, vol. 51, no. 3, pp. 385–463, 2004.

[42] Stefan Kaczmarz, "Angenäherte auflösung von systemen linearer gleichungen," *Bulletin International de lAcademie Polonaise des Sciences et des Lettres*, vol. 35, pp. 355–357, 1937.

[43] Thomas Strohmer and Roman Vershynin, "A randomized kaczmarz algorithm with exponential convergence," *Journal of Fourier Analysis and Applications*, vol. 15, no. 2, pp. 262–278, 2009.

[44] Caihua Chen, Bingsheng He, Yinyu Ye, and Xiaoming Yuan, "The direct extension of admm for multi-block convex minimization problems is not necessarily convergent," *Mathematical Programming*, vol. 155, no. 1-2, pp. 57–79, 2016.

[45] Ruoyu Sun, Zhi-Quan Luo, and Yinyu Ye, "On the expected convergence of randomly permuted ADMM," *arXiv preprint arXiv:1503.06387*, 2015.

[46] Mingyi Hong and Zhi-Quan Luo, "On the linear convergence of the alternating direction method of multipliers," *arXiv preprint arXiv:1208.3922*, 2012.

[47] T. Lin, S. Ma, and S. Zhang, "On the convergence rate of multi-block ADMM," *arXiv preprint arXiv:1408.4265*, 2014.

[48] X. Cai, D. Han, and X. Yuan, "The direct extension of ADMM for three-block separable convex minimization models is convergent when one function is strongly convex," *Optimization Online*, 2014.

[49] Daniel A Spielman and Shang-Hua Teng, "Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems," in *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*. ACM, 2004, pp. 81–90.

[50] Daniel A Spielman and Shang-Hua Teng, "Nearly linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems," *SIAM Journal on Matrix Analysis and Applications*, vol. 35, no. 3, pp. 835–885, 2014.

[51] Ioannis Koutis, Gary L. Miller, and Richard Peng, "Approaching optimality for solving sdd linear systems," *SIAM Journal on Computing*, vol. 43, no. 1, pp. 337–354, 2014.

[52] Alan Edelman, "Eigenvalues and condition numbers of random matrices," *SIAM Journal on Matrix Analysis and Applications*, vol. 9, no. 4, pp. 543–560, 1988.

[53] David Young, "Iterative methods for solving partial difference equations of elliptic type," *Transactions of the American Mathematical Society*, vol. 76, no. 1, pp. 92–111, 1954.

[54] George E Forsythe and Ernst G Straus, "On best conditioned matrices," *Proceedings of the American Mathematical Society*, vol. 6, no. 3, pp. 340–345, 1955.

[55] Olof B Widlund, "On the effects of scaling of the peaceman-rachford method," *Mathematics of computation*, vol. 25, no. 113, pp. 33–41, 1971.

[56] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM Journal on Imgaging Science*, vol. 2, no. 1, pp. 183–202, 2009.

[57] Yi Yang and Hui Zou, "A coordinate majorization descent algorithm for ℓ_1 penalized learning," *Journal of Statistical Computation and Simulation*, vol. 84, no. 1, pp. 84–95, 2014.

[58] James R Angelos, Carl C Cowen, and Sivaram K Narayan, "Triangular truncation and finding the norm of a hadamard multiplier," *Linear algebra and its applications*, vol. 170, pp. 117–135, 1992.